WEAVE SPRINT
PAPER IV · VOLUME SEVEN

MD and DNB psychiatry, theory revision

Neurology, Medicine and Recent Advances

Populations and the Medical Interface

Psychiatry at the scale of a population: the burden, the treatment gap and the programmes meant to close it. Then psychiatry inside general medicine.

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 Seven: Populations and the Medical Interface

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 last volume widens and then narrows. Community psychiatry asks what the discipline looks like when the unit is a population rather than a patient, which is where the treatment gap stops being a statistic. The medical interface chapter closes the whole set, and its discipline is knowing which questions belong to someone else.

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 X Community Psychiatry and Public Health 11

- National programmes and mental health legislation 11
- Epidemiology, burden and treatment gap 48
- Suicide prevention and farmer's suicide 66
- Stigma and anti-stigma interventions 90
- Psychiatric rehabilitation and community care 107
- Primary care, integration and service delivery models 135
- Special-setting and disaster mental health 159
- Prevention, economics and digital mental health 183
- Culture, global mental health and traditional healing 210
- Consolidate and apply 221

PART XI Medical and Systemic-Disease Psychiatry 236

- Cardiac, respiratory and critical-illness psychiatry 236
- Renal, transplant and perioperative psychiatry 270
- Oncology, palliative and paraneoplastic psychiatry 283
- Autoimmune and inflammatory systemic disease, and burn injury 297
- Consolidate and apply 325

BACK MATTER

Subject index 341

About the authors 344

PART X

X

Community Psychiatry and Public Health

Psychiatry at the scale of a population: the burden, the treatment gap, the programmes meant to close it, and prevention that is more than advice.

01 National programmes and mental health legislation

02 Epidemiology, burden and treatment gap

03 Suicide prevention and farmer's suicide

04 Stigma and anti-stigma interventions

05 Psychiatric rehabilitation and community care

06 Primary care, integration and service delivery models

07 Special-setting and disaster mental health

08 Prevention, economics and digital mental health

09 Culture, global mental health and traditional healing

10 Consolidate and apply

PART 10 • Community Psychiatry and Public Health

VISUAL 10.1

India's public mental-health stack



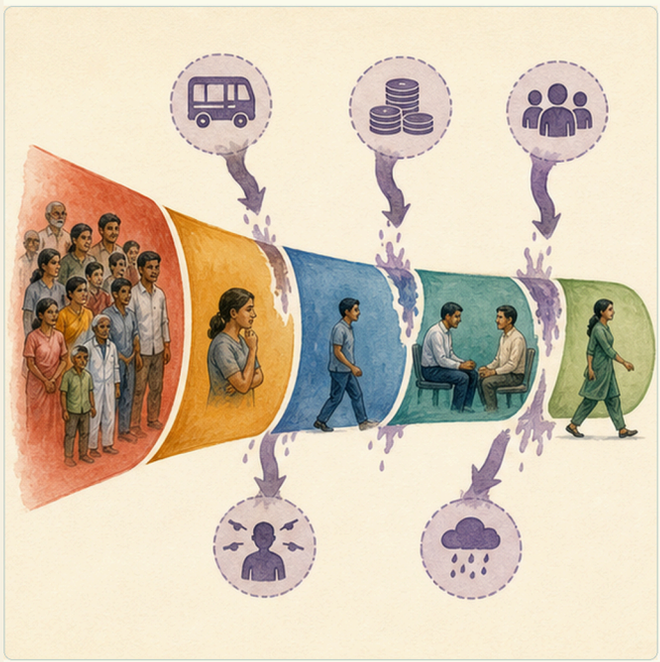
- National policy
- District services
- Primary care
- Tele-mental-health
- Rights and legal protection
- Community and family

Public mental-health care depends on linked policy, district implementation, primary care, tele-services, rights protection and community support.

Source lock: day-39.final.html#sec-t09

VISUAL 10.2

The treatment gap is a leaking pathway

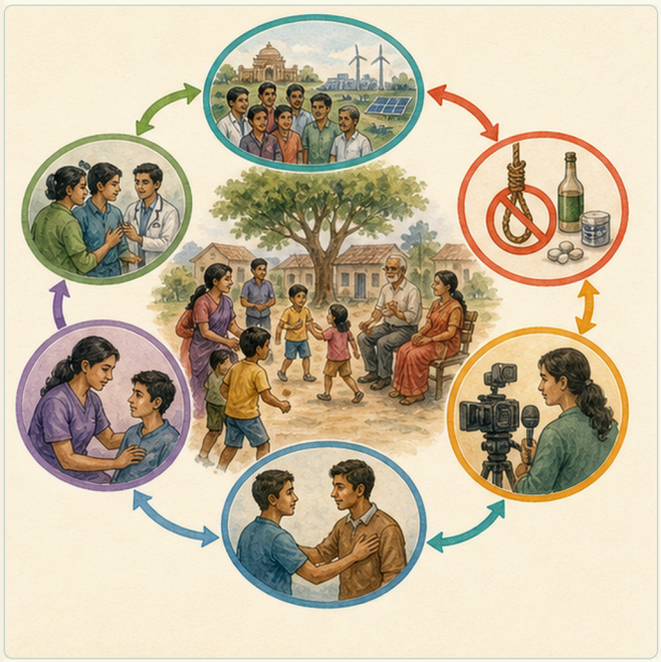


- Need
- Recognition
- Access
- Treatment
- Continuity
- Structural barriers

The treatment gap accumulates across recognition, access, initiation and continuity, shaped by structural barriers rather than a single point of failure.

Source lock: day-39.final.html#sec-t11

VISUAL 10.3



Suicide prevention works across multiple layers

- Population policy
- Means restriction
- Responsible media
- Gatekeepers
- Urgent clinical care
- Aftercare

Suicide prevention combines population measures, safer environments, responsible communication, gatekeepers, timely clinical care and aftercare.

Source lock: day-39.final.html#sec-t16

VISUAL 10.4



Rehabilitation connects treatment to citizenship

- Person-defined goals
- Family and community
- Skills
- Education and employment
- Clinical care
- Participation and recovery

Psychiatric rehabilitation links clinical care to skills, relationships, work, participation and person-defined recovery.

Source lock: day-39.final.html#sec-t29

PAPER IV

Community psychiatry and public health

Nine bands . one complete notes document

National programmes and mental health legislation

1 · Community psychiatry and public health: what this chapter owns, and what it points to

A programme that exists on paper is not a service, and a service that exists in a district is not yet care that a patient has received. The distance between those three things is the subject of this chapter. The national programmes, the money that moves them, the routes by which a person with an illness arrives at somebody able to treat it, and the evidence that any of it worked are all statements about a system rather than about a patient. That change of unit is easy to describe and hard to hold under examination pressure, because every other chapter in this paper has trained you to answer at the level of the person sitting in front of you.

So it is worth stating plainly what this chapter is for: programmes, financing, access pathways, service architecture, population interventions and their evaluation. It is not a second course in mental health law, not a course in professional ethics, not a course in bedside consultation, and not a course in epidemiological method, although it borrows vocabulary from all four and is examined alongside them. The marks here are given for saying what a health system does, at which level it does it, with what resources, for which population, and how anybody would know whether it made a difference. A candidate who answers that with the provisions of an Act, or with the assessment of one patient, has written a good answer to a question nobody asked.

1.1 The lens: the programme, the service and the population

The organising move is a change of unit. Clinical psychiatry takes the individual patient as the unit of analysis and asks what is wrong and what should be done about it. **Public-health implementation** takes the population and the service as its unit and asks a different set of questions: how many people in a defined group carry a condition, how many of them reach treatment, where along the route the remainder are lost, who is expected to act at each level of the health system, and what evidence exists that the arrangement changes anything at all. The clinical facts do not disappear when the lens changes. They stop being the answer and become inputs to a question about delivery.

Two constructs carry most of the weight and recur in every band ahead. **Service architecture** is the arrangement of who does what at which level: the community health worker, the medical officer in a primary care facility, the district team, the general hospital psychiatry unit, the state mental health institution and the specialist tertiary centre. **An access pathway** is the route an individual actually travels, from the first recognition that something is wrong to the point where an effective treatment is delivered, including the long stretches of that route that lie outside the

health system entirely. Almost every intervention taught in these notes acts on the architecture or on the pathway, and deciding which is usually the first useful thing to say about it.

RULE **The answer is written at the level of a population or a service, not at the level of one consultation.** Before answering any question in this territory, say which population is being described and which level of the system is expected to act. A clinically impeccable account of how a patient should be assessed and treated earns very little here, because the question was about who reaches that assessment, who is available to deliver it, and who pays for it. Keep the clinical content in a clause and spend the answer on the delivery decision.

1.2 What this chapter deliberately does not teach

These notes border four others closely enough that material leaks in both directions, and each border is worth rehearsing as a sentence. Naming the owning chapter accurately is faster and better rewarded than rebuilding its material from memory, and it is the only way an examiner can see that you knew where the boundary was.

- The statutes themselves are not taught here. Their provisions, the machinery they create, the rights they confer, the safeguards they impose and the procedures they prescribe belong to the Mental health legislation and forensic psychiatry notes. What belongs here is the service consequence: what a law obliges a system to provide, who has to be in post before that provision can exist, and how it is financed and counted.
- Professional conduct, capacity and consent, custodial procedure and the statutory frame governing disaster response belong to the Forensic ethics, consent and legal procedure notes. These notes take the same events from the service side: how a response is organised, who staffs it, at which phase, and what happens to routine care while it runs.
- The bedside act belongs elsewhere as well. The cultural formulation of an individual case, the idioms in which a particular patient describes distress, the named culture-bound presentations and the clinical assessment of a patient already under the care of a traditional or faith healer are taught in the Consultation-liaison, geriatric and transcultural psychiatry notes. This chapter takes the same ground at population level: help-seeking, the sequence of providers a community actually consults, and the relationship between the health system and the other places people go.
- Method belongs to the Research Methods, Biostatistics and Epidemiology notes: what incidence and prevalence are as measures, how a survey is designed and sampled, what a study design can and cannot support, and how a rate is standardised. These notes consume the products of that method rather than teaching it, and the discipline they add is about reading a published figure correctly rather than about producing one.

Two narrower handovers are worth the same treatment. The clinical assessment and management of the individual patient at risk of suicide belongs to the Somatic-symptom, sexual, impulse-control disorders and suicide notes, while the population-level prevention of suicide is taught here in full. The clinical care and protective assessment of the individual child belongs to the

Child psychiatry: assessment, treatment, services and protection notes, while the service and programme interface with child protection law sits here.

TRAP **Answering a programme question by reciting the statute.** Asked how mental health services are organised in a district, candidates produce the sections and safeguards of the governing legislation, because that material is adjacent enough to feel responsive and was revised harder. The recitation may be accurate and still earn little, because the examiner asked what is delivered, by whom and at what level. State the legal obligation in a clause, then spend the answer on the structure built to discharge it.

1.3 What makes a national figure true

This is the most number-heavy territory in the paper, and almost every figure in it is one you have heard quoted at a seminar. That is exactly the problem. A national figure carries meaning only when four things travel with it, and a figure quoted without them is not a fact but a rumour with a decimal point.

Source	Reference year	Population	Denominator
WHICH BODY ISSUED IT	THE PERIOD IT DESCRIBES	WHOM IT COUNTS	WHAT IT IS A PROPORTION OF

Each field fails in its own way. Drop the source and the number rests on the speaker, which under examination means on memory. Drop **the reference year** and a survey that describes one period is offered as a description of the present, which matters more here than anywhere else in the paper because the national instruments in this field are re-run rarely and their findings stay in circulation long after the period they measured. Drop the population and the commonest error in the whole territory follows: the same survey's finding for one group is quoted as though it applied to another, so a figure that is perfectly accurate becomes a false statement about the wrong people. Drop **the denominator** and a count is read as a rate.

The last field needs one remark, because the honest answer is sometimes that no valid denominator exists. A total of sanctioned units, a cumulative count of calls to a helpline and a national count of deaths in an occupational group are all real numbers with no meaningful proportion attached, either because nobody counted the base or because the base is not knowable. The correct sentence says so, and stops: a cumulative count of calls is a count of calls, never a count of people. Refusing to supply a denominator is not weakness in an answer. It is the only version of the sentence that is true.

PITFALL

Reading a sanctioned figure as a delivered one. Programme documents in this field count what has been approved, funded or covered on paper, and those counts are published, quotable and entirely genuine. They are not counts of functioning services. A district may appear in an approved list with posts unfilled, and a facility may be listed with no specialist attached. The same trap runs through activity data: a service that reports contacts is reporting contacts, not distinct individuals, and certainly not individuals who completed treatment. Before quoting any coverage figure, ask what the noun beside the number actually is, and say that noun out loud in your answer.

1.4 The map of the territory

Nine bands follow. They are set out below not as a list to memorise but as a routing table: given a question, it says which band holds the answer. The bands themselves supply the programmes, the figures, the models and the attributions.

BAND	WHAT IT ESTABLISHES	THE KIND OF ANSWER IT REWARDS
National programmes and the legislative frame around them	What the state has undertaken to provide, through which programme structures, and at which level of the system each acts	Components and objectives, the level each one operates at, and the difference between what is sanctioned and what is running
Burden and the treatment gap	How much illness a defined population carries, and what proportion of it reaches treatment	A figure with its source, its reference period, its population and its denominator attached, and the determinants behind the shortfall
Suicide prevention at population level	Which population-level actions bear on deaths by suicide, and which groups carry raised risk	Actions classified by the population they address rather than by the individual they treat, with the level of action named
Stigma and the interventions aimed at it	How stigma operates in the public, in the person and in institutions, and what has been tried against each	The type of stigma identified first, then the intervention matched to that type and the evidence behind it
Rehabilitation and community care	What happens after the acute episode, and where a person then lives, works and is supported	The components of a rehabilitation package and the setting each belongs to, with residential and daytime options distinguished
Primary care, integration and delivery models	Who delivers mental health care where no specialist exists, and what supports them in doing it	The task each cadre performs, the supervision structure behind it, and the model that organises both
Special settings and disaster	How services are organised for populations that ordinary outpatient care does not reach	The constraints the setting imposes first, then the service response adapted to those constraints
Prevention, economics and digital care	What can be prevented, what illness costs, how care is paid for, and what technology changes about reach	The level of prevention named, the economic question stated precisely, and a claim about reach kept separate from a claim about benefit
Culture, global mental health and traditional healing	How explanatory models and existing sources of help shape the route into care	The pathway described as it is actually travelled, and the working relationship between the health system and other providers of help

1.5 The spine underneath the map

The bands are not nine separate topics; they are one argument told in sequence. Burden establishes the size of the problem in a defined population. The treatment gap converts that burden into a statement about the service, because a gap is not a fact about illness at all but a measure of how badly access has failed. The pathway work then explains where along the route people are lost, and the architecture says who might have caught them at each of those points. The intervention bands, from suicide prevention through stigma, rehabilitation and primary care integration, are all attempts on one or the other: they change who can act, or they change how a person reaches somebody who can. Prevention and economics ask which of those attempts is worth making and who pays for it. Evaluation asks whether any of it moved.

Read that way, the last band is not an appendix. Culture, explanatory models and existing healers sit at the beginning of the pathway rather than at the end of the chapter, because for a large part of the population the first consultation for a mental health problem is not with the health system at all. A reader who meets that band last but places it first has understood the sequence.

1.6 Inputs, outputs and outcomes

Evaluation is the part of this territory candidates most often skip, and it is where an understood answer separates fastest from a memorised one. Programme evidence comes in three kinds, and confusing them is the single most productive source of wrong statements in the whole field.

WHAT THE FIGURE COUNTS	A QUESTION IT CAN ANSWER	A QUESTION IT CANNOT ANSWER
An input: what was committed, sanctioned, staffed or funded	Whether the system undertook to act, and on what scale	Whether anything was delivered to anybody
An output: what the service did, counted as activity	Whether activity took place, and how much of it	Whether the people who most needed the service were the ones who got it
An outcome: what changed in the population	Whether the health of a defined group moved, and in which direction	Which part of the arrangement produced the change, unless the evaluation was designed to attribute it

The distinction has a companion in the idea of **functional coverage**, which asks not how many units exist but how many are staffed, supplied, open and used by the population they were built for. Sanctioned coverage is an input; functional coverage sits between an input and an output; and neither is an outcome. Where the functional figure has not been measured, the correct statement gives the approved figure, says that it is an approval, and declares the working figure unavailable. That declaration is a complete answer, and a better one than a confident number that quietly promotes a sanction into a service.

GOOD TO REMEMBER

Before quoting any national figure in a case presentation, a ward round, a teaching session or an answer script, ask four questions in order. Who issued this number. What period does it describe, as distinct from when the document was published. Which population does it count, and is that the population I am about to apply it to. And what is it a proportion of, or is it a bare count with no valid denominator. A clinician who runs those four before speaking will not be the person who quotes a survey's finding for one group as though it described another, which is the most frequently repeated error in this whole field.

1.7 What a good answer in this territory looks like

A small number of moves earn most of the marks here, and they recur across all nine bands in different clothing. They are worth learning as habits, because the detail here changes while the moves do not.

- Naming the level of the health system at which the answer operates, before describing what happens there. Community, primary care, district, general hospital, state institution and tertiary centre are not interchangeable settings, and most questions here are really asking which one you would use.
- Separating what a policy or programme undertakes from what a service delivers, and saying which of the two the evidence you are quoting actually measured.
- Attaching the source, the period, the population and the denominator to every figure, and declining to give a figure at all where any of the four is unavailable to you.
- Handing over what a neighbouring chapter owns in a single sentence that names it, rather than reconstructing statute, ethics, bedside method or epidemiological technique from memory in the middle of a service answer.

Those four are also, in the same order, the four ways an answer here goes wrong. An answer written without a level reads as an opinion about mental health in general. A figure without its scope is a wrong number wearing a citation. And an answer that drifts into the law has left the marks on the table while writing accurately about something else.

1.8 A scaffold to carry into essays and vivas

Under examination pressure the detail goes first, and in this field the detail is exactly what is hardest to hold, because it consists of programme names, structures, reference periods and figures that resemble one another. What has to survive is a routing sequence. The scaffold below is a teaching device rather than a validated framework, and it will structure an answer on any band in these notes, from a programme to a delivery model to a question about culture and help-seeking.

MNEMONIC

Population, Pathway, Programme, Proof

Population: name the group the question is about and, if a number is coming, what it is a proportion of. **Pathway**: describe how that group actually reaches care and at which step it stops, including the steps taken outside the health system. **Programme**: name the level of the system expected to act, what has been put there to act with, and who supervises it. **Proof**: say what would count as evidence that it worked, and whether the figure you have is an input, an output or an outcome.

For a short answer, run the four and stop. For a long essay, expand each into a paragraph and take the specifics from the band that owns them. In a viva, make the sequence audible: say whom you are talking about, how they reach care, who is meant to be there, and how you would know whether it helped. That last sentence separates a candidate who has learned a list of programmes from one who has understood why they exist, and it is also what a district team needs from the psychiatrist it consults.

A programme is not a service, a service is not access, and access is not care: name which of the three a question is asking about, and attach the source, the period, the population and the denominator to every figure you offer.

2 · National Mental Health Programme (NMHP): components, objectives and evolution

A national programme is not a policy, not a statute and not a scheme, and the marks in this area are lost by candidates who cannot tell the four apart. The **National Mental Health Programme** is the instrument through which the Union government made mental health a delivery obligation of the general health system rather than the business of a handful of specialist hospitals. Everything else in Indian community psychiatry hangs off it: the district limb that carries services to the population, the training and workforce apparatus, the money that reaches a state, and the reporting through which any of it becomes visible. Understand what kind of object it is and the rest of this chapter becomes a set of consequences.

This section teaches the programme as programme architecture: what it is for, what it consists of, how it has moved, how it is financed and monitored, and what can honestly be said about how far it has reached. Its district implementation is a separate section of this chapter. Its role as the delivery frame for community rehabilitation of a person with a psychotic illness, set against assertive community treatment, is taught in the Schizophrenia: management, antipsychotics and adverse effects notes, and is not restated here. The statutes the programme now runs inside are taught in the Mental health legislation and forensic psychiatry notes.

2.1 Four instruments, four different kinds of authority

Before the programme, the vocabulary. Indian mental health administration uses four words that examiners deliberately mix, and each carries a different kind of force.

INSTRUMENT	WHAT IT FIXES	WHO IS BOUND	HOW FAILURE SHOWS
Statute	Rights, duties, offices, procedures and remedies	Everyone within the jurisdiction, on pain of legal consequence	An entitlement exists and is not honoured; the remedy is legal
Policy	Direction, priorities and the populations named as vulnerable	Government, as a statement of intent	The direction is stated and nothing is funded to move in it
Programme	Objectives, components, delivery architecture and reporting	The health system that operates it	Approved on paper, unstaffed in practice
Scheme	A defined benefit or activity with its own budget line	The implementing agency and the beneficiary	Money released, benefit not received

The programme sits in the third row, and its characteristic failure mode is the third one. That single row is the most examinable idea in this section and it recurs in every part of this chapter: an approved unit is not a functioning unit, a sanctioned post is not a filled post, and a released instalment is not a service delivered. Hold the distinction and you will not misread a coverage figure again.

RULE **A programme is an architecture plus a reporting system, and it can only be judged on what it reports.** Its objectives tell you what it was designed to change; its components tell you the machinery it intends to use; its monitoring tells you what evidence of change will ever exist. Where the third is thin, no amount of the first two will let you say whether the programme worked.

2.2 Objectives, and the one this chapter can put a source against

The objectives of this programme are among the most recited facts in Indian psychiatry and among the least reliably sourced. The founding programme document, the one that sets out the original objectives, could not be retrieved from a serving official source during the preparation of this chapter: the ministry's own page returned nothing after a domain migration, a repair attempt returned the same dead address, and no archived snapshot was available. That is stated here rather than papered over, because the honest position is itself an answer. The objectives as they are usually recited in teaching may well be accurate; this chapter simply cannot certify them against a document it could open, and it will not print a list it cannot stand behind.

What can be certified is the objective the programme's current official page carries. The Directorate General of Health Services states an objective of the National Mental Health Programme in these terms: **to enhance human resource in mental health sub-specialties**. That is a live official statement, read from the programme's own live page, and it describes the mental health subspecialty workforce of India. It is not a rate and it has no denominator, which is correct and not a defect: an objective is a direction of travel, not a proportion of anything.

Notice that this objective is itself evidence of evolution, which is the second half of what this item covers. An objective about human resource appearing on the current page tells you that the

programme, whatever it was designed to do at the outset, now regards the production of specialist and subspecialist manpower as one of its own deliverables rather than as something the education system supplies to it from outside.

IN INDIA

Because the programme names enhancement of human resource in mental health sub-specialties as an objective of its own, a department planning against it should treat training capacity as a programme output to be planned and reported, not as an academic by-product of clinical work. That changes what a unit counts: not only patients seen, but trainees trained, supervisors available, and the subspecialty posts that a trained person can actually move into. A workforce objective without a post at the end of it produces qualified people for other systems.

2.3 Why human resource is the binding constraint on this programme

Public health programmes usually fail on one of four constraints: money, materials, people or information. Mental health services are unusual in that the first two are comparatively cheap. The essential medicines are old, off patent and inexpensive; the physical infrastructure is a room, a record and a reliable supply. What the service cannot substitute for is trained people, because the intervention *is* a trained person's time, whether that person is a psychiatrist, a clinical psychologist, a psychiatric social worker, a nurse or a trained medical officer at the first level of care.

That has three design consequences which run through the rest of this chapter.

- **Capacity is built slowly and lost quickly.** Specialist training takes years, while a transfer, a resignation or an unfilled sanctioned post removes the capacity immediately. A programme that expands its coverage faster than it produces people converts a staffing problem into a coverage illusion.
- **Task-shifting is not optional, it is arithmetic.** If the specialist workforce cannot be enlarged at the rate the population needs, some part of the work must move to a broader cadre with supervision. The mechanism, the guidance behind it and the medical officer's role in it are taught in their own sections of this chapter.
- **Retention is a programme variable.** Where a trained person is posted, whether the post is at the level they were trained for, and whether supervision exists, determine whether training converts into service. None of that is visible in a count of trainees.

2.4 The architecture, from the national programme down to the district

A national programme in India is not delivered nationally; it is delivered by states, in districts. Health is administered at state level, so the Union instrument sets objectives, defines components, provides a share of the money and asks for returns, while the actual service is run by a state health department through a district unit. That two-level structure explains most of what is puzzling about Indian programme data. The centre knows what it approved and what it released; the district knows what it staffed and what it delivered; and the two are not the same fact.

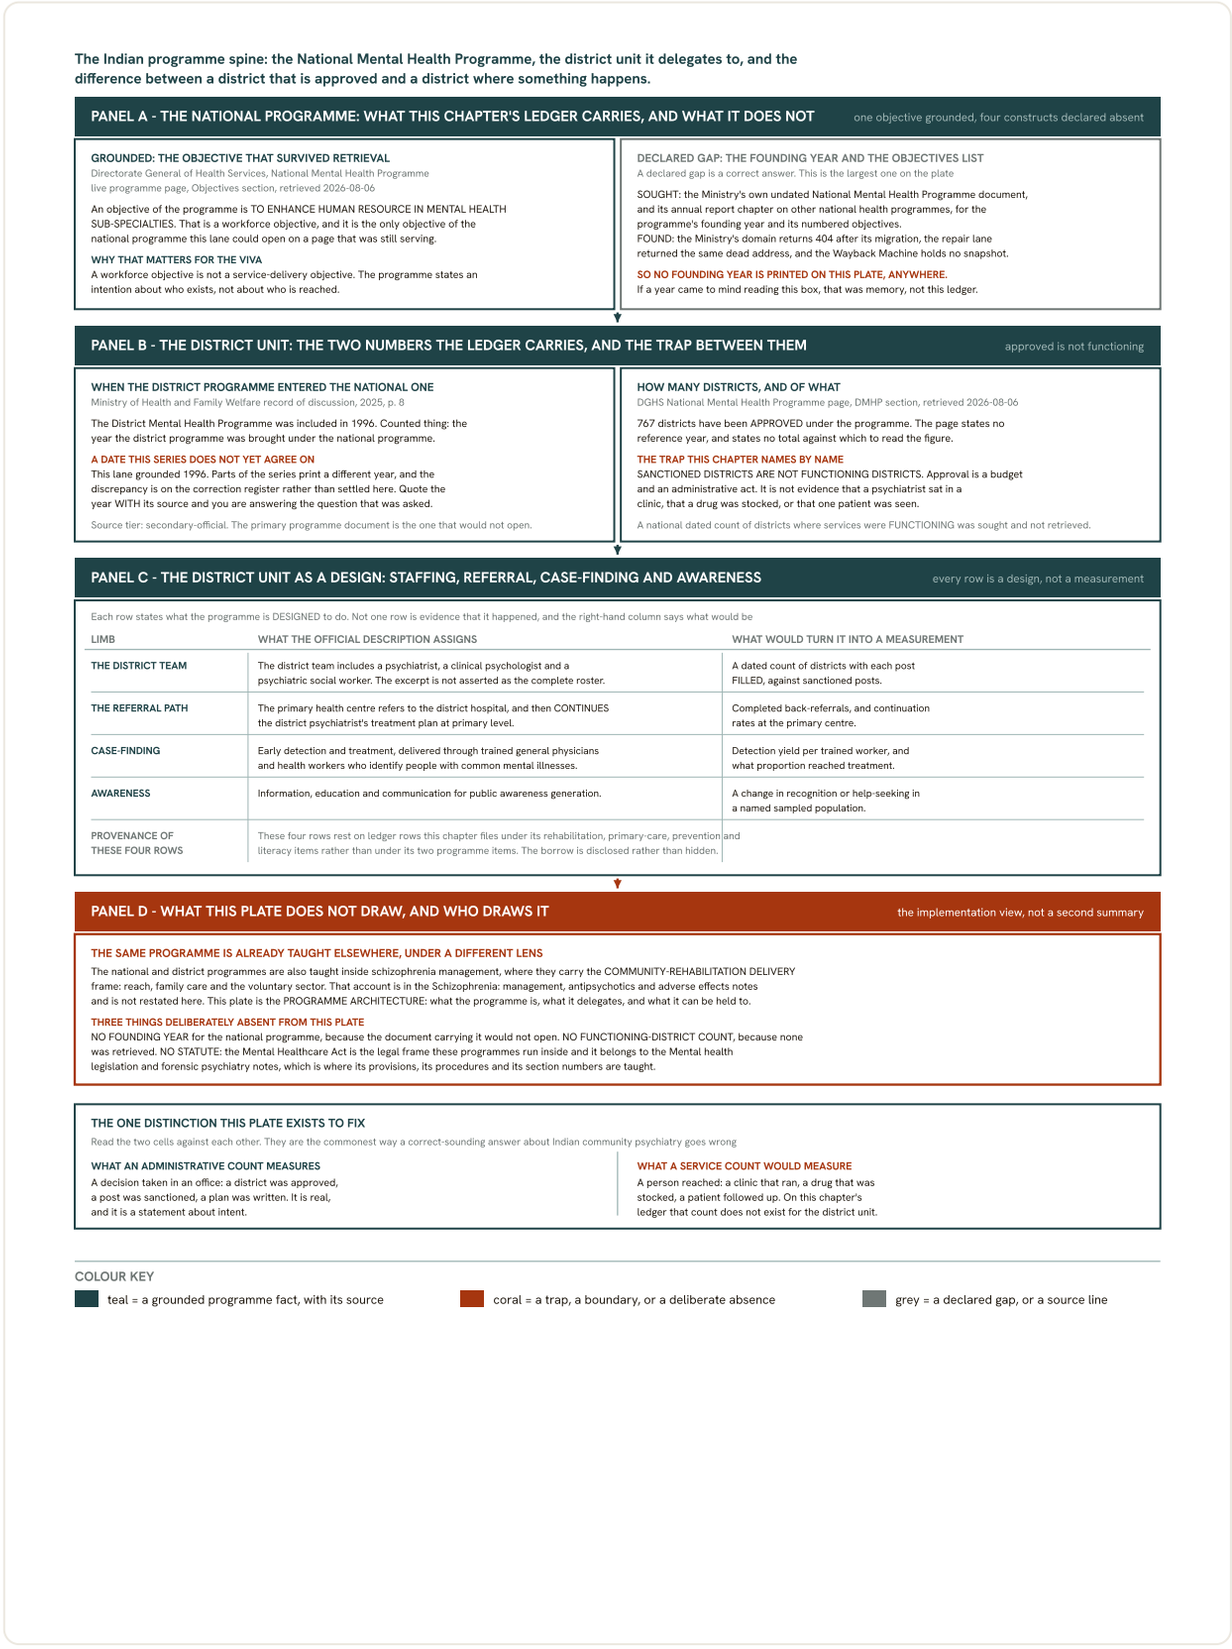


Fig 39.1 **The Indian programme spine, drawn as an implementation view.** The national programme contributes one retrievable objective, to enhance human resource in mental health sub-specialties; its founding year and numbered objectives are a declared gap, because the Ministry's own documents no longer serve and no archived snapshot exists, so no year is printed. The district programme was included in 1996 and is reported as approved in 767 districts, with no reference year and no stated total. The chapter's own trap sits between those two facts: approval is an administrative act, and a national count of districts where services were functioning was sought and not retrieved. The staffing, referral, case-finding and awareness rows state design, never delivery.

The accompanying figure sets out that spine, from the national programme down through the district unit to the components, staffing and referral arrangements at the point of delivery, and marks the boundary where the community rehabilitation account in the Schizophrenia: management, antipsychotics and adverse effects notes takes over. The district limb, the District Mental Health Programme, has its own section immediately after this one and its structure, components, staffing and coverage evidence are taught there rather than here.

What belongs here is the logic of the relationship. A national programme can do four things and only four: it can set an objective, define a component, move money, and require a return. Everything else, including whether a post is filled, whether a drug is on the shelf and whether a patient is seen, happens at a level the national instrument can influence only through those four levers. When you are asked why a national programme underperforms, the answer is almost always that one of the four levers is weak, and most often it is the fourth.

2.5 Evolution, and the document that could not be opened

The programme has moved considerably since it began: from an initial concentration on integrating mental health into general health care, through the addition of a district implementation limb, towards a wider portfolio that now includes workforce production, tertiary capacity and digital service delivery. The programme's own annual reporting is where those movements are recorded plan period by plan period, and it is the document a candidate should ideally cite.

Two of this chapter's declared evidence gaps sit exactly there. The ministry annual report chapter that records the programme's position among the other national health programmes could not be retrieved from a serving official address: the domain returned nothing after migration, a repair pass returned the same dead address, and no archived copy was available. Two separate records in this chapter's evidence ledger point at that same unreachable section. The consequence is that this chapter describes the direction of the programme's evolution, which is visible from its current official page and from the instruments that surround it, and prints no dated milestone list, because the document that would date the milestones could not be opened.

TRAP **Writing the year a programme started from memory.** Programme years, policy years and Act years are the most quoted and most misquoted facts in this whole area, and they feel exactly as certain from the inside as a year you have actually read in a document. This chapter prints a programme year only where a source it could open carries one. Where it could not, it names the programme without the year and says so. In a viva, "the founding document is not currently retrievable from an official source, so I would not quote a year I cannot cite" is a stronger answer than a confident wrong date.

2.6 How you would actually judge this programme, and what a figure must carry

Judging a national programme means choosing a denominator, and most published commentary skips that step. There are four honest questions and they get progressively harder to answer.

- **Reach.** How many districts, units or facilities are covered? This is the easiest to answer and the easiest to misread, because approval and function are recorded by different people for different purposes.
- **Function.** Of the covered units, how many are staffed, supplied and open? This is the number that matters and the number most often missing.
- **Contact.** Of the people in the population who need care, how many make contact with the service? This requires a population estimate as a denominator, which comes from a survey rather than from the programme.
- **Outcome.** Of those who make contact, how many are adequately treated and how many improve? This requires a follow-up system, which is the weakest layer of Indian programme reporting.

Source	Year	Population	Denominator
WHICH BODY ISSUED THE FIGURE	WHICH YEAR IT REFERS TO, NOT WHEN PUBLISHED	WHICH PEOPLE IT DESCRIBES	WHAT IT IS A PROPORTION OF

Those four tiles are the discipline this entire chapter runs on, and they are why so few programme figures appear in it. A national number missing any one of them is not a fact; it is a recollection with a decimal point attached. Where a programme figure carries all four, this chapter prints it with all four. Where it does not, the chapter names the construct without its number and records the absence, which is the correct answer and not an evasion.

MNEMONIC

Objectives, components, delivery, money, returns

The five things any national health programme consists of, in the order you should describe one. An answer that gives the objectives and stops has described a fifth of the instrument.

PITFALL

Describing the programme entirely through its objectives. Objectives are the most memorable part and the least informative, because every national programme's objectives read approximately alike. What differentiates one programme from another is its delivery architecture, the level at which the money stops, and what it requires anyone to report. If you can describe those three for the National Mental Health Programme, you can answer almost any question in this area, including ones about programmes you have never studied.

The National Mental Health Programme is a Union instrument that sets objectives, defines components, moves a share of the money and requires returns; states and districts do the delivering, which is why approval and function are never the same number.

3 · District Mental Health Programme (DMHP): structure, components and implementation

The district is where Indian mental health policy either becomes a service or stays a sentence. The **District Mental Health Programme** is the operational limb through which the National Mental Health Programme reaches a population: a team, a budget line, a place to work from, a referral route in both directions, and a training obligation towards the general health workforce of the district. Every national ambition in this chapter, integration into primary care, task-shifting, closing the treatment gap, reaching a rural population, is discharged at this level or nowhere.

This section teaches the district limb as a service architecture and as an implementation problem. The parent programme's objectives, evolution and national machinery are taught in the section immediately before this one and are not repeated here. The way this same district structure functions as the community delivery vehicle for rehabilitation in schizophrenia, set beside assertive community treatment, is taught in the Schizophrenia: management, antipsychotics and adverse effects notes.

3.1 Why the district, and not the state or the block

The choice of the district as the unit is not administrative convenience; it is the level at which the three things a mental health service needs coincide. The district is small enough that a single specialist team can physically reach its facilities and large enough to justify one existing at all. It is the level at which the health administration already holds a budget, a stores system, a supervisory hierarchy and a reporting line, so a new programme can be attached to machinery that exists rather than built beside it. And it is the level at which secondary care sits, which matters because a mental health service that has no bed and no acute pathway becomes an outpatient clinic that refers its sickest patients away.

Compare the alternatives and the logic becomes obvious. A state-level unit can plan but cannot deliver; the distances defeat it and the specialists sit in the capital. A block-level unit can reach people but cannot be staffed with a specialist team at that density, and it has no bed. The district is the compromise, and like every compromise it inherits both problems in miniature: it is still far from the periphery, and it is still thin at the top.

RULE **A district programme is defined by what it does to the district's existing services, not by what it does itself.** Its clinic is the smallest part of it. Its real product is a general health system that can identify, start, continue and refer mental health care, which is why training and supervision, not consultations, are the components that determine whether it works.

3.2 When the district limb came in, and what that means

The **District Mental Health Programme** was brought in under the national programme in **1996**, and that year is stated in an official Ministry of Health and Family Welfare consultation record rather than in the founding instrument itself. That distinction matters and is worth carrying: the source is official but secondary, meaning it restates the fact rather than being the document that created it. A different year for the same event circulates widely in teaching

material, and this chapter prints the year its own retrievable official source carries rather than the year that is most often repeated.

The substantive point behind the date is that the district limb was an addition. The national programme began as an integration argument directed at the general health system; the district structure came afterwards as the vehicle for making that argument operational. That sequence explains a persistent feature of Indian mental health administration, which is that the district programme is always described as a component of something larger and is always funded as one, so its fortunes rise and fall with decisions taken about the parent programme rather than about itself.

3.3 What a functioning district unit actually consists of

Describe a district unit by its components and you will be able to say, for any real district, exactly which part has failed. Six things have to be simultaneously true.

- **A team.** A specialist core with psychiatry, clinical psychology, psychiatric social work and psychiatric nursing, plus programme and record staff. The composition matters less than the completeness: a team missing its psychosocial cadre becomes a prescribing clinic, and a team missing its record staff becomes invisible to the reporting system.
- **A place, at more than one level.** A base at the district hospital and outreach to peripheral facilities, because a service that exists only at the district headquarters has simply relocated the travel burden from the state capital to the district town.
- **An uninterrupted drug supply.** The essential psychotropics are cheap and off patent, so a stock-out is a logistics failure and never a cost one. A single interrupted supply undoes months of adherence work and teaches a family that the service cannot be relied on.
- **Referral in both directions.** Upward referral is easy to arrange and easy to overuse. Back-referral, the deliberate return of a stabilised patient to a peripheral facility with a plan the medical officer can follow, is the harder half and the one that determines whether the specialist team ever has capacity for new patients.
- **A training cascade.** The district team trains medical officers, health workers, counsellors and teachers, who then identify and manage what they can. This is the multiplier; without it the programme's reach is capped at the number of patients four or five specialists can personally see.
- **Records and returns.** Whatever is not recorded did not happen as far as the programme is concerned, and whatever is recorded without a denominator cannot be interpreted by anyone above the district.

Read that list again as a diagnostic instrument. When a district programme is described as non-functional, one of those six has failed, and the failure is usually the team or the supply. When it is described as functional but ineffective, the failure is almost always the fifth or the sixth.

3.4 The training cascade, and the school component that belongs elsewhere

The cascade is the part of the design that carries the public health logic. A specialist team that spends its time in direct consultation converts a scarce resource into a small number of individual encounters. The same team spending part of its time training, supervising and supporting a much larger general workforce converts the scarce resource into system capacity. That is the argument for task-shifting expressed at district scale, and the guidance and the medical officer's role in it are taught in their own sections of this chapter.

The programme's school outreach component is part of the same cascade, and it is taught in the Child psychiatry: assessment, treatment, services and protection notes rather than here. That chapter carries the school component in detail, including the teacher-training model with its own coverage figures for master trainers and schools. This section names the component as one limb of the district cascade and reproduces none of its numbers, so that a reader who meets a teacher-training figure meets it once, in one place, with one set of denominators.

3.5 Coverage: the number that exists, and exactly what it counts

There is one national coverage figure for this programme that can be printed with its provenance intact, and reading it correctly is the single most examinable skill in this section. The Directorate General of Health Services reports that **767 districts** have been approved under the programme. That is approved coverage. It is not functioning coverage, and the source itself does not state the total number of Indian districts against which it should be read, so even the proportion approved cannot be computed from it without importing a denominator from somewhere else.

767	1996	Not stated
DISTRICTS APPROVED UNDER THE PROGRAMME, NOT DISTRICTS FUNCTIONING	YEAR THE DISTRICT LIMB CAME IN UNDER THE NATIONAL PROGRAMME	DISTRICT TOTAL, SO NO APPROVAL PROPORTION IS COMPUTABLE

Set the four fields against that figure and its limits are plain. The source is the programme's own official page. The population is districts of India considered for approval. The counted thing is districts approved. The denominator is not supplied by the source at all. Three of the four are strong; the fourth is missing, and the honest sentence says so rather than filling the hole with a district total remembered from elsewhere.

3.6 The number that does not exist, and why that is the answer

The obvious next question is how many of those districts have a service actually running in them. That figure was sought for this chapter from the national programme, from the directorate and from the ministry, including by targeted searches of official domains for functioning or operational districts, and by attempted retrieval of the ministry's own annual report. It was not obtained. There is, in this chapter's evidence position, no national, reference-dated, official count of districts in which district mental health services were functioning.

That absence is content, not a hole in the teaching. A candidate who says that the approved figure is the only national coverage number available, that approval is a sanctioning decision rather than

a service fact, and that the functional count is not published in a citable form has described the Indian evidence position accurately. A candidate who converts the approved figure into a claim about services running has made an error that no amount of confidence repairs.

COVERAGE QUESTION	WHAT WOULD ANSWER IT	WHAT EXISTS NATIONALLY
Is the district sanctioned?	An approval count from the programme's own records	Published: 767 districts approved under the programme
Is the service running?	A reference-dated count of districts with services functioning	Sought from the programme, the directorate and the ministry; not obtained
Is the team complete?	Filled posts against sanctioned posts, by cadre and by district	Not held in this chapter's evidence
Are people being reached?	Contacts against an estimated population needing care	Requires a survey denominator, taught in the survey section of this chapter

TRAP **Reading a sanctioned district as a functioning district.** This is the classic error in Indian community psychiatry and it is built into the way the data are produced: approval is recorded centrally because the centre approves, while function is a state and district fact that nobody aggregates. The approved count is real, official and correctly reported; it simply answers a different question from the one the reader is usually asking. Say "approved" every time you say the number.

IN INDIA

When a district is described to you as covered by the programme, treat that as an administrative statement and ask three follow-up questions before you plan anything on it: which posts in the team are filled today, whether the essential psychotropics were in stock last month, and whether any patient has been back-referred to a peripheral facility with a written plan. The published national figure is 767 districts approved under the programme, and approval is where those three questions begin rather than where they end.

3.7 What monitoring would have to look like

The reason the functional count does not exist is not indifference; it is that nobody has defined the unit of observation. Deciding whether a district service is functioning requires an operational definition that a clerk can apply consistently: how many days open, which cadres present, what stock position, what minimum activity, over what reference period. Without that definition, every state answers a slightly different question and the national total means nothing.

- An operational definition of a functioning unit, applied uniformly, with a stated reference period.
- Filled against sanctioned posts by cadre, because a team is not staffed if only the psychiatrist post is filled.
- Stock-out days for the essential psychotropics, which is the fastest single indicator of whether a service is real.

- Back-referrals as a proportion of stabilised patients, which measures whether the cascade is working rather than whether the clinic is busy.

MNEMONIC

Team, place, drugs, referral, training, returns

The six components of a district unit, and the six ways one fails. Name the six and you can diagnose any district service you are asked about, including one you have never visited.

PITFALL

Judging a district programme by its outpatient numbers. A busy district clinic can coexist with a completely failed cascade, and often does, because direct consultation is the easiest activity to increase and the easiest to count. Rising attendance at the district base with no back-referral and no trained medical officers means the specialist team has absorbed the district's workload instead of distributing it, and the service will saturate.

The published national figure is 767 districts approved under the programme; approval is a sanctioning decision, no national count of functioning districts is published, and the two must never be quoted as one.

4 · Mental Healthcare Act 2017: key provisions, advance directives, nominated representative

Every programme in this chapter runs inside a statute, and the statute is not this chapter's to teach. The Mental Healthcare Act is the legal frame within which India's national and district mental health programmes now operate. It fixes the entitlements a service has to honour, it creates instruments a patient may use to direct their own future care, and it establishes bodies that supervise both. Its provisions, including the advance directive and the office of the nominated representative, are taught in full in the Mental health legislation and forensic psychiatry notes. This section names the frame, prints that pointer, and stops.

That is a deliberate structural decision and not an omission. A statutory entitlement is a text with a fixed wording, and the fastest way to corrupt it is to let two chapters paraphrase it independently. The second account will differ from the first in a modal verb, a threshold or a procedural step, and neither reader will be able to tell which drifted. One owner, one wording.

4.1 The boundary, drawn once, and applied everywhere in this chapter

The division of labour is easier to hold if it is stated as a question rather than as a list of topics. A statute answers **what a service owes**: the right, the duty, the instrument, the procedure and the remedy. A programme answers **what a service has**: the posts, the drugs, the budget line, the referral route and the coverage. The two questions have different owners because they have different failure modes. A wrong entitlement misleads a patient about their rights; a wrong coverage figure misleads a planner about a population.

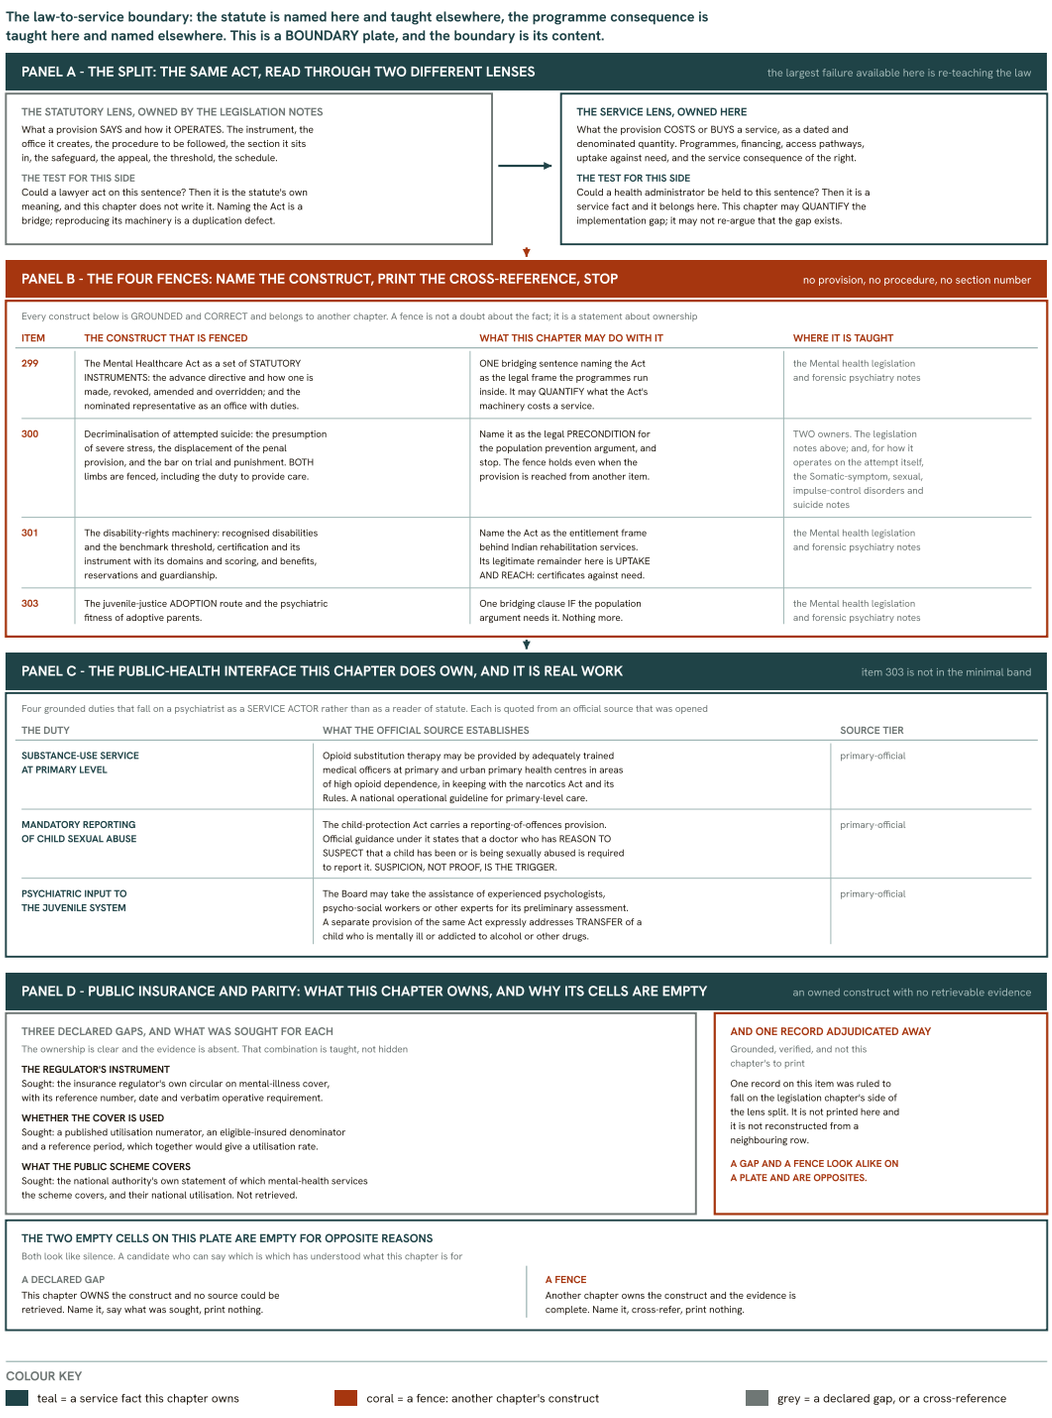


Fig 39.2 The law-to-service boundary, drawn as an exclusion map. The statute is named here and taught in the Mental health legislation and forensic psychiatry notes; what a provision costs or buys a service is taught here. Four constructs are fenced: the Act's statutory instruments, the decriminalisation of attempted suicide, the disability-rights machinery, and the juvenile-justice adoption route. No provision, procedure or section number is reproduced on this plate. The interface this chapter does own is real work and it is grounded: opioid substitution therapy provided by trained medical officers at primary level, a doctor's duty to report suspected child sexual abuse on reasonable suspicion rather than proof, expert psychological and psycho-social assistance to the juvenile-justice Board, and a transfer route for a child who is mentally ill or addicted. The insurance and public-scheme cells are empty because the regulator's own instrument could not be retrieved, which is the opposite kind of emptiness from a fence.

The accompanying figure is the exclusion map this chapter applies whenever a statute appears in it. Everything on the legal side is named here and taught elsewhere. Everything on the service side, meaning programmes, financing, access pathways, service architecture and population consequences, is taught here in full. No arrow crosses the line without a pointer attached to it.

RULE **Naming a statute is not teaching it, and pointing at its owner is the assertion that you did not teach it.** A section that names the Act and gives no provision, no procedure and no locator has discharged its obligation exactly. A section that adds one more paragraph has started writing a second copy of somebody else's chapter.

4.2 What legitimately remains here, and why it is not written

There is real work on this side of the line, and it is a counting job rather than an interpretive one. Statutory machinery has a running cost, and that cost is a service fact. An entitlement to make an advance directive implies somewhere it is lodged, somebody who retrieves it at the point of admission, and a record system that can produce it under pressure. An office such as the nominated representative implies people who know it exists, staff trained to identify one, and documentation that survives a transfer between institutions. Supervisory bodies imply members, sittings, quorum and a caseload. None of that is law; all of it is establishment, workload and coverage, which is what a public health chapter measures.

- Sites at which the Act's instruments can actually be lodged, retrieved and acted on, counted against the sites that admit patients.
- Staff trained to operate the machinery, counted against the establishment a district service is sanctioned for.
- Supervisory bodies constituted and sitting, counted against the units they are meant to supervise, with the reference year attached.
- The interval between an entitlement taking legal effect and the service being resourced to deliver it.

Each of those is a legitimate sentence for this chapter, and each becomes a defect the moment it is written without a source, a reference year, a population and a denominator. This chapter's evidence ledger holds no such dated, denominated quantity for the implementation load of this Act. The consequence is stated plainly rather than covered over: the argument that the Act's machinery is unevenly resourced is already in print in the legislation chapter and is not repeated here, and the quantification that would be this chapter's own contribution to it does not yet exist in a form that can be cited.

TRAP **Answering a service question with a statutory one.** Asked what the Act changed for a district service, candidates recite the rights. The examiner has already heard the rights. What distinguishes an answer is the operational consequence: who now has to be found, what now has to be recorded, and what has to be staffed for the entitlement to be real. If you cannot put a number on that, say so and say why, because at present nobody can.

Name	Point	Stop
THE ACT, AS THE LEGAL FRAME A SERVICE RUNS INSIDE	TO THE CHAPTER THAT TEACHES ITS PROVISIONS	NO PROVISION, NO PROCEDURE, NO LOCATOR HERE

MNEMONIC

Name it, point at it, price it

The three moves of a boundary section. The first two are always available. The third is the only one that belongs to this chapter, and it is the one that needs evidence this chapter does not hold for this Act.

The Mental Healthcare Act is named here as the legal frame the programmes operate inside; its provisions, instruments and procedures are taught in the Mental health legislation and forensic psychiatry notes, and this section reproduces none of them.

5 · Rights-based provisions and decriminalisation of suicide under the MHCA 2017

One provision of the Mental Healthcare Act changed the legal status of the suicide attempt in India, and it is the precondition for everything this chapter says about population suicide prevention. Before it, a person who survived an attempt met the health system through a door that also led to a police station. After it, the attempt is treated in law as the presentation of a person under severe stress rather than as an offence to be tried. That change is the reason a national prevention argument can be made at all in this country, because a population strategy cannot ask people to come forward while the coming forward is itself an offence.

The provision itself, its wording, the presumption it creates, the offence it displaces and the duty it places on government are taught in the Mental health legislation and forensic psychiatry notes, and the way the same provision operates on the individual suicide attempt is taught in the Somatic-symptom, sexual, impulse-control disorders and suicide notes. Both limbs are named here and neither is reproduced. This section is the bridge between them and the programmes that follow.

5.1 Why a public health chapter names a legal change and then stops

The rights architecture of the Act and the **decriminalisation** of the attempt are single texts with fixed wordings, and a statutory entitlement paraphrased twice in one book is an entitlement that will eventually be printed at two strengths. The risk is not abstract. The difference between a duty and a discretion lives entirely in the modal verb, and a public health author reaching for a summary sentence is precisely the author most likely to soften one. So the boundary here is drawn hard: no wording, no procedural step, no locator, no penalty, and no restatement of the duty that the appropriate government carries. Those belong to the two chapters named above and are complete there.

RULE **The fence on this construct holds wherever you approach it from.** The same provision is reached from the suicide prevention strategy, from psychiatric rehabilitation and from the emergency presentation, and it is fenced identically each time. A reader who meets it in three different sections should meet the same pointer three times, not three summaries that have drifted apart.

5.2 The service consequence, which is what remains here

Decriminalisation is a legal fact with an epidemiological shadow, and the shadow is this chapter's business. Three service consequences follow from it, and none of them requires reproducing a line of the statute.

- **Presentation.** A legal deterrent to disclosure suppresses presentation, and suppressed presentation is invisible in service data. When the deterrent is removed, recorded contacts can rise while the true rate is unchanged or falling. A service that reads a rise in attendances after a legal change as a rise in incidence has misread its own denominator.
- **Recording.** National mortality and morbidity data on suicide in India come through a police-recorded route, and what a legal category counts is not what a clinical category counts. The counting problem is taught with the national figures in the sections on suicide epidemiology and on farmer suicide in this chapter.
- **Obligation.** The provision does not merely remove a penalty; it places an affirmative duty of care on the state. A duty of that kind is discharged by services that exist, are staffed and are reachable, which is a coverage question, and coverage is what the programme sections of this chapter measure.

The third of those is where a candidate can distinguish themselves and where this chapter is most exposed. Naming the duty is the legislation chapter's work. Quantifying its discharge, meaning dated and denominated counts of the services that would actually deliver it, is this chapter's work, and it is done in the sections that own the national and district programmes rather than here.

TRAP **Answering "what did decriminalisation change" with the law rather than with the service.** The legal answer is already owned and already examined elsewhere. What a public health answer adds is the mechanism by which a legal change alters who arrives, what gets recorded and what the state is now obliged to provide. Candidates who recite the provision have answered a different question competently.

Named	Owned	Kept
THE LEGAL PRECONDITION FOR POPULATION PREVENTION	ELSEWHERE, IN BOTH ITS LEGAL LIMBS	PRESENTATION, RECORDING AND THE DUTY TO PROVIDE

MNEMONIC

Who arrives, what is written down, what is owed

The three service consequences of a legal change, and the only three this section carries. Everything else about the provision is taught in the two chapters named here.

Decriminalisation of the suicide attempt is named here as the legal precondition for population suicide prevention in India; its wording, its presumption, the offence it displaces and the duty it creates are taught in the Mental health legislation and forensic psychiatry notes and, as they operate on the attempt itself, in the Somatic-symptom, sexual, impulse-control disorders and suicide notes.

6 · Rights of Persons with Disabilities Act 2016 and mental illness as disability

Rehabilitation in India is not only a clinical activity; it is an entitlement system, and the Rights of Persons with Disabilities Act is the instrument that creates the entitlements. Everything this chapter later says about supported living, supported employment, day care, family-based care and long-term community support sits downstream of a legal architecture that recognises mental illness as a disability and attaches benefits to that recognition. This section names that architecture and points to where it is taught.

The Act itself, the conditions it recognises, the threshold at which a disability becomes a benchmark disability, the certification process, the Indian Disability Evaluation and Assessment Scale used within it, and the benefits, reservations and guardianship arrangements that follow, are taught in full in the Mental health legislation and forensic psychiatry notes. No schedule, no percentage, no procedure, no instrument domain and no score range is reproduced here. The whole band belongs there and it is complete there.

6.1 The one distinction worth carrying away from this boundary

An entitlement has two separate existences and candidates routinely collapse them. There is the **statutory entitlement**, which is what the law says a person is owed, and there is **uptake**, which is the proportion of the people who qualify who actually hold the document that unlocks it. The first is fixed the day the statute commences. The second is a service and administrative outcome that can stay near zero for years afterwards while the entitlement remains, on paper, complete.

That distinction is the reason a rights statute appears in a public health chapter at all. Legislation converts a clinical population into an eligible population. Certification converts an eligible population into a benefiting population. The gap between those two conversions is a coverage problem with all the ordinary features of one: it is unevenly distributed, it rewards the literate and the connected, it costs the applicant time and money, and it is invisible in any dataset that counts only the people who succeeded.

RULE **A right is a legal fact; a certificate is a service outcome; a benefit received is a population outcome. They are three different numbers and only the first is guaranteed by the statute.** An answer that moves from the Act straight to the benefit has skipped the step where Indian implementation actually fails.

6.2 What remains here, and the shape of the missing evidence

The legitimate remainder on this item is reach, and reach is a ratio. A defensible sentence about disability certification for mental illness in India would name an issuing authority, a

reference year, a numerator of certificates issued for mental illness, and a denominator of people estimated to meet the threshold. It would then say what proportion of the eligible population holds the document, and how that proportion varies by state, by sex and between rural and urban residence.

- Certificates issued for mental illness, nationally, in a stated year, from a named issuing authority.
- An estimate of the population meeting the threshold, derived from a survey whose own reference year and population are stated.
- The ratio of the two, presented as coverage and never as prevalence.
- The interval between application and certificate, and the proportion of applications that fail, because a process people abandon has the same population effect as a process that refuses them.

This chapter's evidence ledger holds none of those quantities for this item. That is stated here in prose rather than hidden in a footnote, because it is examinable in itself: a resident who can say that the entitlement is statutory, that the binding step is certification, and that the national uptake ratio is not established in a citable form has given a better answer than one who quotes a benchmark percentage from a slide and attaches it to the wrong denominator.

TRAP **Quoting a disability threshold as though it were a prevalence, or as though holding it were automatic.** The threshold is a legal criterion applied to an individual assessment; it says nothing about how many people cross it and nothing about how many of those who cross it are certified. The threshold, the assessment instrument and the certification procedure are all taught in the Mental health legislation and forensic psychiatry notes and none of them is reproduced in this chapter.

Named	Owned	Kept
THE ACT, AS THE ENTITLEMENT FRAME FOR REHABILITATION	SCHEDULE, THRESHOLD, INSTRUMENT AND BENEFITS, ELSEWHERE	UPTAKE AND REACH, WHICH ARE NOT YET GROUNDED

MNEMONIC

Eligible, certified, benefiting

Three shrinking populations, one statute. The law fixes the first, the certification process fixes the second, and only the third is a service outcome anybody experiences. This chapter owns the shrinkage; the statute belongs elsewhere.

The Rights of Persons with Disabilities Act is named here as the entitlement frame behind Indian rehabilitation and community services, and is taught in the Mental health legislation and forensic psychiatry notes; what belongs to this chapter is uptake, meaning certificates issued against estimated need, and no such national ratio is available to be printed.

7 · National Mental Health Policy and Tele-MANAS initiative

A policy states a direction and a programme builds something; India acquired one of each within a decade, and the pair is the cleanest worked example of the difference in this whole chapter. The National Mental Health Policy set out what the country said it wanted mental health care to become. **Tele-MANAS**, the national tele mental health programme, is a piece of national infrastructure that was actually built, that runs continuously, and that publishes reach figures. Teaching them together lets you see exactly where a stated intention turns into a countable service, and exactly where the counting starts to mislead.

Two internal boundaries apply throughout. The clinical workflow of a remote consultation and the regulatory framework governing it are taught in this chapter's section on telepsychiatry, and evidence on apps and digital products is taught in its section on digital mental health. This section owns the policy instrument and the national programme architecture, and nothing else.

7.1 The Policy: what a policy instrument is for

India's first national mental health policy was launched in **October 2014**. Being first matters more than the date itself. Until then the country had a programme without a policy, which is an odd arrangement: an implementation machine running without a published statement of what it was implementing towards. A policy supplies that statement. It names the values, the priorities, the populations regarded as needing particular attention, and the sectors expected to contribute.

What a policy cannot do is compel. It creates no right, imposes no duty and releases no money. Its influence is exerted through three indirect routes: it gives programme designers something to point at when they ask for resources, it gives advocates a published standard against which to measure performance, and it sets the vocabulary in which subsequent instruments are written. Judged as a legal instrument a policy is weak; judged as an agenda-setting instrument it can be decisive.

RULE **Ask of any policy what it would take for it to have been implemented, and whether anybody is required to report on that.** A policy with no named indicator and no reporting obligation cannot be shown to have failed, which sounds like protection and is actually the reason it can be quietly abandoned.

7.2 The scope the Policy claimed, and why that scope is unusually wide

The most consequential sentence in the Policy is not about services at all. Its stated scope includes **addressing the social determinants of mental health, including poverty, environmental issues and education**, as reproduced in the official national survey report's account of the mental health system. A policy that commits itself to social determinants has committed to something no health department can deliver alone, and that is deliberate rather than careless.

Three consequences follow, and they are examinable.

- **Intersectoral by construction.** If poverty and education are inside the scope, then departments of rural development, social welfare, labour and school education are

implementing agencies whether or not they know it. The health department becomes a coordinator as well as a provider.

- **The indicator problem becomes hard.** Service indicators are countable: contacts, admissions, prescriptions. Social determinant indicators belong to other sectors, move slowly, and cannot be attributed to a mental health policy even when they improve.
- **The accountability question sharpens.** A wide scope with no cross-sectoral reporting obligation produces a policy that is easy to agree with and impossible to audit. That is the single most useful criticism a candidate can make of any national mental health policy, and it applies well beyond India.

Which populations the Policy identifies as needing particular attention, and one population that a published policy analysis found it does not name among them, are examined in the Perinatal/women's mental health and emergency psychiatry notes. That finding is presented there as an analysis of the instrument rather than as a statement made by the instrument, and this chapter does not re-argue it.

7.3 Tele-MANAS: infrastructure rather than intention

The Government of India launched the national tele mental health programme, **Tele-MANAS**, on **10 October 2022**. Understand first what problem it is shaped around. India's specialist workforce is small and geographically concentrated, and it cannot be enlarged quickly. A telephone service breaks the geographical constraint immediately: it decouples where the caller is from where the responder is, it operates continuously rather than during clinic hours, and it can be staffed from wherever trained people happen to live.

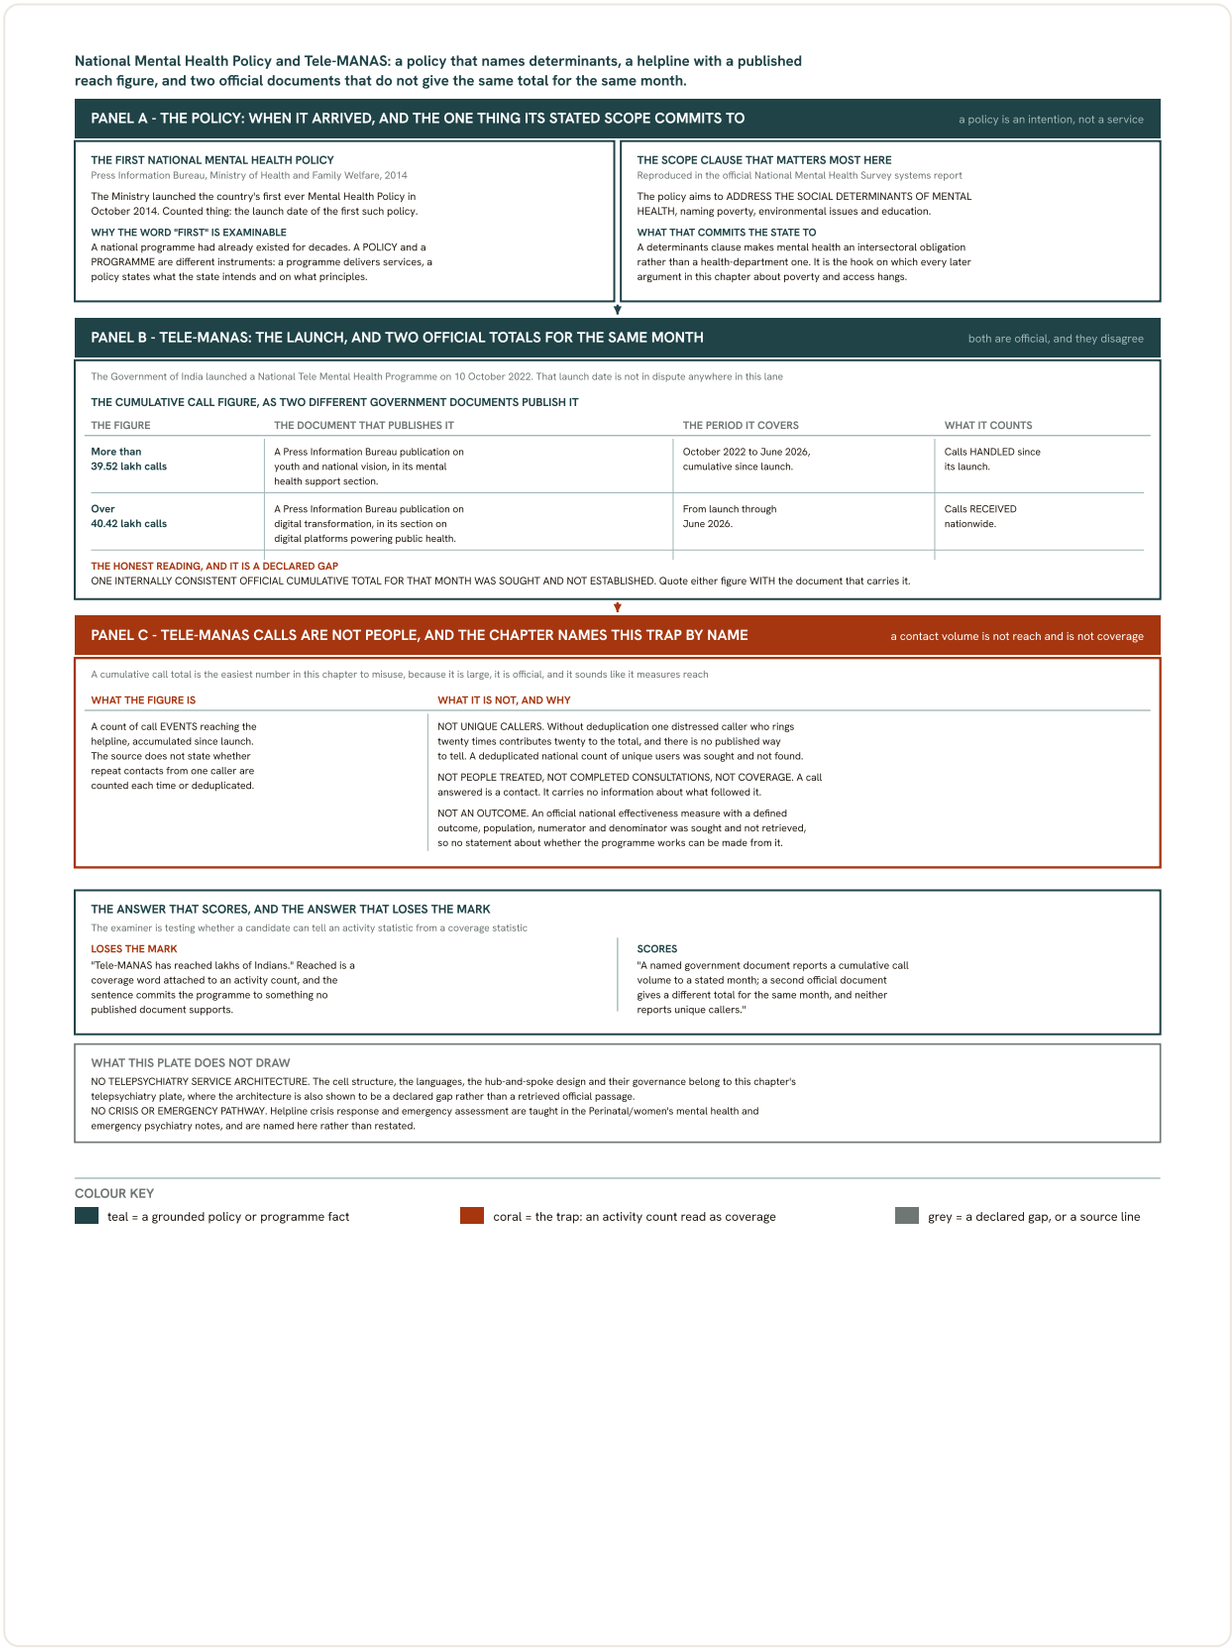


Fig 39.3 Policy and Tele-MANAS: an intention, a helpline, and a number that two official documents report differently. The country's first national Mental Health Policy was launched in October 2014 and its stated scope commits to addressing the social determinants of mental health, naming poverty, environmental issues and education. The national tele mental health programme was launched on 10 October 2022. For the same month one government publication reports more than 39.52 lakh calls handled since launch and another reports over 40.42 lakh calls received; one internally consistent official total was sought and could not be established, so either figure must be quoted with the document that carries it. Neither is a count of people: no deduplicated national count of unique users exists and no official effectiveness measure with a stated numerator and denominator was retrieved.

The accompanying figure sets out the policy instrument and the programme architecture together, including the tiered structure by which a call is answered, escalated and routed onwards, and it sets out alongside them exactly what the published reach figures do and do not count. Read the architecture as three separable claims: that a call is answered, that it is answered by somebody competent, and that it connects to onward care where onward care is needed. Only the first of the three is directly visible in any published figure.

7.4 The reach figures, and the disagreement between two official documents

This is the passage to read twice, because it is the chapter's whole discipline compressed into one construct. Two official documents report the cumulative volume this service has handled, and they do not agree.

One official release reports that Tele-MANAS handled **more than 39.52 lakh calls** by June 2026, counted cumulatively since launch, the reference period running from October 2022 to June 2026. A separate official release reports that the service had received **over 40.42 lakh calls** nationwide as of June 2026, again counted from launch. Two arms of the same official information apparatus, the same construct, the same stated month, two different totals.

10 Oct 2022	39.52 lakh	40.42 lakh
TELE-MANAS LAUNCH	CALLS HANDLED SINCE LAUNCH, ONE OFFICIAL RELEASE, TO JUNE 2026	CALLS RECEIVED SINCE LAUNCH, A SECOND OFFICIAL RELEASE, SAME MONTH

Do not reconcile them. There is no internally consistent official cumulative total for that month available to this chapter; the programme's own live dashboard was sought and could not be reached. The correct answer names both, attributes each to its own release, notes that one speaks of calls handled and the other of calls received, and says that the difference between those two verbs is not defined in either document. A resident who can do that has understood something more valuable than a memorised figure: that two official numbers for one construct usually differ because they are counting slightly different events, and that the difference is invisible unless you read the noun each one attaches to.

7.5 What a call total is not

Both figures count call events. Neither counts people. No official deduplicated national count of unique users of this service since launch was obtained for this chapter, and it was sought from the ministry and from the programme's own dashboard. That is not a small omission; it is the difference between a reach statistic and a workload statistic.

QUESTION ABOUT A NATIONAL HELPLINE	WHAT IT NEEDS	EVIDENCE POSITION IN THIS CHAPTER
How much work does it do?	A cumulative count of call events, with a reference period	Two official cumulative totals, differing, both printed above with their sources
How many people has it reached?	A deduplicated count of unique callers	Sought from the ministry and the programme dashboard; not obtained
What proportion of those who need it use it?	Unique callers over a population estimate of need	Not computable, because the numerator above does not exist
Does it help?	An operationally defined outcome, with numerator, denominator and reference period	Sought; no official national effectiveness measure was obtained

Work down that table and the shape of the evidence becomes clear. The service publishes the number that is easiest to produce and hardest to misinterpret innocently. Each successive row is more useful to a planner and less available. That pattern is not peculiar to this programme; it is what happens whenever a service is monitored by the system that runs it rather than by one that evaluates it.

TRAP **Converting calls into people.** A cumulative call total counts contacts, and neither source states whether repeat contacts from the same caller are counted each time or removed. One distressed person calling weekly for a year contributes many calls and one user. Writing that the service has "reached" that many people, or dividing the total by the population to produce a coverage rate, manufactures a statistic that no document supports. Say "calls handled, cumulative since launch, as reported by" and then name the release.

IN INDIA

When you are asked to justify a tele mental health service in a proposal or a viva, quote the volume as volume and then say what you would need to justify it as reach: a deduplicated user count, a population denominator and a defined outcome. Both published totals here are cumulative call counts, more than 39.52 lakh in one release and over 40.42 lakh in another for the same month, and neither is a count of people. Building a service case on the larger figure without that qualification is how a planning document acquires an error nobody can later trace.

7.6 Effectiveness, and the measure that is not published

The question that matters clinically is whether calling the service leaves people better off. Answering it requires an operationally defined outcome, a stated population, a numerator and a denominator, and a reference period. No official national effectiveness measure of that kind was obtained for this chapter, despite searches of the programme dashboard and ministry statements.

It is worth being precise about what that absence does and does not mean. It does not mean the service is ineffective; nothing here supports that claim either. It means the evidence that would settle the question in the direction of either answer is not published in a citable form, and that any confident statement about the programme's effect, positive or negative, is currently an

opinion. For an examination, the sequence to give is: what is published, what is not, and what would be needed. That sequence works for every programme in this chapter.

MNEMONIC

Volume, users, coverage, outcome

Four questions about any national service, in increasing order of usefulness and decreasing order of availability. Tele-MANAS publishes the first. The second, third and fourth were sought for this chapter and not obtained.

PITFALL

Treating a helpline as a substitute for a service rather than as a front door to one. A call answered competently is a triage and a beginning; its value depends almost entirely on whether the person can be connected to continuing care afterwards, which is a district-level question. A national telephone layer built over a thin district layer moves the bottleneck rather than removing it, and it makes the bottleneck harder to see, because the call is counted and the failure to connect onwards is not.

The Policy states a direction and creates no duty; the national tele mental health programme is real infrastructure whose published totals count calls, cumulatively since launch, and never people.

8 · Narcotic Drugs and Psychotropic Substances Act, Protection of Children from Sexual Offences Act and Juvenile Justice Act as they intersect psychiatry

Three statutes reach into psychiatric practice from outside psychiatry, and each one hands a mental health service a population it did not choose and a duty it did not write. The narcotic drugs legislation, the NDPS Act, defines who is an offender and who is a patient in substance use. The child sexual offences legislation, POCSO, places an affirmative reporting duty on the doctor who suspects abuse. The juvenile justice legislation creates forums that need psychiatric expertise and provides for children in institutional care who turn out to be mentally ill or dependent on substances. None of the three was written by health administrators, and all three determine what a health service has to be able to do.

This section teaches the public health and service interface only: prevention, referral, the population consequences of a legal frame, and the intersectoral machinery each statute presumes. The statutes themselves, their provisions, classifications, procedures, penalties and time frames, are taught in the Mental health legislation and forensic psychiatry notes, and the child protection procedure and the juvenile justice machinery are taught in the Child psychiatry: assessment, treatment, services and protection notes. No provision is reproduced here.

8.1 One question asked of three statutes

The productive way into this material is not to summarise three Acts, which is what an unstructured answer does and what the owning chapters already do properly. It is to ask each

statute the same two questions and compare the answers.

STATUTE	THE POPULATION IT CONSTRUCTS	WHAT IT ASKS OF A MENTAL HEALTH SERVICE
Narcotic drugs legislation	People who use controlled substances, split between an offender frame and a treatment frame	Treatment capacity that a person can reach without the reaching being risky, and services able to operate lawfully within the drug-control rules
Child sexual offences legislation	Children who have been or are being sexually abused, and the adults who encounter them	A clinician who recognises and reports, and a service organised so that reporting leads somewhere
Juvenile justice legislation	Children in conflict with law and children needing care and protection	Expert assessment capacity for its forums, and a route for children in institutional care who are mentally ill or substance dependent

Read across the third column and a single service requirement appears three times: someone competent has to be reachable, and there has to be somewhere for the person to go afterwards. That is a district service question, and it is why these three statutes belong in a public health chapter at all rather than only in a legal one.

RULE **A statutory duty on a clinician is only as real as the service standing behind it.** Legislation can require a doctor to act; it cannot by itself create the trained person, the pathway or the receiving service that makes the action useful. The gap between the duty and the capacity is a health system problem, and it is the part of this material that belongs here.

8.2 Drug control as a public health instrument

Every drug-control regime does two things at once, and the two pull in opposite directions for a health service. Supply reduction attacks availability through prohibition, interdiction and penalty. Demand reduction attacks consumption through prevention, treatment and rehabilitation. A statute that carries both creates a population that is simultaneously a target of enforcement and a target of care, and the person in that population has to decide which of the two they are approaching when they walk into a clinic.

That ambiguity has measurable service consequences and they are the substance of this subtopic.

- **Delayed presentation.** Where disclosure carries legal risk, people present later, sicker, and often through an emergency rather than through a clinic. The delay is invisible in service data because the earlier presentations that never happened leave no trace.
- **Incomplete disclosure.** A person who fears legal consequence under-reports quantity, route and the other substances used. Every clinical decision downstream, including dose, is then made on a corrupted history.

- **Family and employer concealment.** Legal risk attaches to the household as well as the individual, so the informants a clinician relies on have their own reason to minimise.
- **Distorted surveillance.** Enforcement data count seizures and cases; health data count contacts. Neither counts prevalence, and a rise in either can reflect activity rather than epidemiology.

The reformative side of the drug-control frame, meaning the statutory routes by which a person who undergoes treatment is treated differently by the law, and the arrangements for recognising treatment institutions, are taught in the Mental health legislation and forensic psychiatry notes. What that chapter also already establishes is that such routes work only where a qualifying institution is genuinely within reach, so this section does not re-derive that point and instead takes it as given.

8.3 Where the treatment actually has to sit

If reach is the binding constraint, then the design question is what level of the health system can lawfully and safely provide substance use treatment. National operational guidance for mental, neurological and substance use care at the first level of care takes a clear position on the hardest case. It states that **opioid substitution therapy may be provided by adequately trained medical officers at primary health centres and urban primary health centres, in areas with high levels of opioid dependence, in keeping with the drug control Act and Rules**.

Read the qualifiers, because they carry the whole design. The provider is a medical officer, not a specialist, which is task-shifting applied to the most tightly regulated treatment in psychiatry. The condition is adequate training, which converts the sentence from a permission into a training obligation. The targeting is geographical, directed at areas where opioid dependence is concentrated rather than spread thinly everywhere. And the whole permission sits inside the drug-control rules, so the regulatory frame is not bypassed but worked within.

IN INDIA

Where opioid dependence is concentrated, do not design the service around a specialist centre that people travel to. National operational guidance places this treatment at the first level of care, with adequately trained medical officers at primary health centres and urban primary health centres as the providers, within the drug-control rules. The practical decisions that follow are about training, supervision and regulated supply at that level, not about building another specialist clinic; a treatment that requires daily or near-daily attendance fails on distance long before it fails on efficacy.

8.4 Mandatory reporting as a service obligation, not only a personal one

The child sexual offences legislation is the point at which a clinician's private judgement stops being private. Official guidance issued under that legislation is explicit about the clinician's position: **a doctor who has reason to suspect that a child has been or is being sexually abused is required to report it**. Two features of that sentence do the work. The trigger is suspicion with a reason behind it, not certainty and not proof, which is a deliberately low

threshold. And the verb is one of requirement, not of discretion, so the decision the clinician faces is not whether to report but how to do it well.

The procedural detail, meaning who the report goes to, what form it takes, what protection the reporter has and what follows for the child, is taught in the Child psychiatry: assessment, treatment, services and protection notes. What belongs here is the systems consequence, which is routinely underestimated.

- **The duty falls on every clinician, so the service must assume every clinician will meet it.** A duty discharged well by the one consultant who happens to know the pathway is a service that has not implemented anything.
- **Confidentiality is reshaped rather than suspended.** The child and the accompanying adult need to be told, in advance and in plain language, what a clinician cannot keep to themselves. Saying it before the disclosure is honest; saying it afterwards is a betrayal the family will remember.
- **Reporting is the beginning of a service obligation, not the end of one.** The child still needs assessment, treatment and follow-up, and the reporting event frequently disrupts exactly the therapeutic relationship in which that would have happened.
- **A service needs a named route and a rehearsed one.** Which staff member acts, whom they contact, what is recorded and who supports the family afterwards should be settled before the first case, because the first case will present out of hours to the least experienced person present.

TRAP **Treating mandatory reporting as an ethical dilemma to be weighed case by case.** Candidates reach for the confidentiality argument because it is the interesting one. Where an official duty to report exists, the weighing has already been done by the legislature and the clinician's judgement is engaged in how to report, how to prepare the child and family, and how to preserve continuity of care afterwards. The examinable skill is the service design, not the deliberation.

8.5 Juvenile justice: two places psychiatry is written into the system

The juvenile justice framework creates forums that a mental health service interacts with directly. The Juvenile Justice Board and the Child Welfare Committee are, from a service point of view, destinations: places a child is routed to and routed from, each with its own decisions to make and its own need for information a clinician can supply. Their constitution, composition and procedure are taught in the Child psychiatry: assessment, treatment, services and protection notes.

Two provisions of that legislation reach psychiatry directly and both are service-shaping rather than merely procedural.

The first concerns expertise. In carrying out its preliminary assessment, **the Board may obtain assistance from experienced psychologists, psycho-social workers or other experts**. The modal is permissive and that is the important part. The statute contemplates expert input rather than compelling it, which means whether the expertise is actually available in a district is a

service question and not a legal one. A permissive provision in a district with no available expert is a provision that does nothing.

The second concerns children already inside the system. The legislation expressly provides for **the transfer of a child who is mentally ill or addicted to alcohol or other drugs**. That is a statutory recognition that institutional care will contain children who need psychiatric or addiction treatment, and it presumes a receiving service exists to transfer them to. Whether it does is, again, a district capacity question.

One further limb of the same legislation, the adoption route and the psychiatric fitness of prospective adoptive parents, is named here in a single clause because a reader tracing the statute will look for it, and it is taught in the Mental health legislation and forensic psychiatry notes. No eligibility criterion, regulation or fitness assessment procedure appears in this chapter.

Reach	Report	Receive
DRUG CONTROL: TREATMENT A PERSON CAN SAFELY APPROACH	CHILD PROTECTION: A DUTY ON EVERY CLINICIAN	JUVENILE JUSTICE: EXPERTISE AND A PLACE TO TRANSFER TO

8.6 What the three have in common, and what a service can do about it

All three statutes make a mental health service a downstream consumer of categories it did not define. The categories are legal, the thresholds are legal, and the timing is set by legal processes. A service that treats this as somebody else's problem discovers its consequences one crisis at a time. A service that plans for it does four things, and none requires a change in the law.

- **Identify the interfaces in advance** and name, for each, the person in the service who acts and the external counterpart they contact.
- **Train for the trigger rather than the statute.** Staff need to recognise the situation that starts the duty; the legal detail can be looked up, and the recognition cannot.
- **Build the receiving capacity, not only the referring pathway.** Every one of these statutes presumes somewhere to send a person. Pathways into services that do not exist generate paperwork and nothing else.
- **Record at the interface.** Cross-sector work is where documentation is weakest and where it is most likely to be examined later.

MNEMONIC

Reach, report, receive

What each of the three statutes asks a mental health service for: treatment a person can safely approach, a clinician who recognises and reports, and a service able to receive a child the system sends. Every one of the three is a capacity question rather than a legal one.

PITFALL

Assuming an intersectoral pathway exists because a statute presumes one. Legislation routinely presumes a counterpart service, a trained expert or a receiving institution, and says nothing about creating any of them. Before promising a court, a committee or a family that a route exists, confirm that the person at the other end of it is currently in post. The commonest failure at these interfaces is not a wrong decision; it is a correct referral into a vacancy.

These three statutes are named here and taught elsewhere; what this chapter owns is the capacity each one presumes, which is treatment a person can safely reach, a clinician who recognises and reports, and a service able to receive the child the system sends.

9 · Ayushman Bharat, insurance parity for mental illness and the IRDAI mandate

Who pays decides who is treated, and in mental health the answer has historically been the household. A person with a chronic psychiatric illness in India typically needs medication continuously for years, review appointments indefinitely, occasional admission, and travel to reach any of it, while earning less because of the illness. Every one of those costs falls on a family whose earning capacity the illness has already reduced. Financing is therefore not an administrative appendix to this chapter; it is one of the determinants of whether care is used at all.

This section teaches the financing and insurance machinery: how publicly financed health insurance is built, what parity means and how it is implemented, why mental illness fits an insurance model badly, and what can and cannot currently be said with evidence about coverage and utilisation in India. The legal right to **insurance parity**, its statutory wording and the account of how that right came in practice to move the insurance regulator, are taught in the Mental health legislation and forensic psychiatry notes. What that chapter does not name is the regulator, and this one does: the Insurance Regulatory and Development Authority of India, the **IRDAI**.

One boundary inside the Ayushman Bharat umbrella needs stating at the outset, because the shared brand causes real confusion. This section teaches the insurance and benefit-design limb, meaning the publicly financed hospitalisation scheme and the regulation of insurance products. The school health and wellness limb of the same umbrella is a different programme with a different delivery mechanism, and it is taught in the Child psychiatry: assessment, treatment, services and protection notes. One brand, two programmes, and their figures must never be pooled.

9.1 Four ways a health system pays, and what each does to psychiatric care

Any health system pays for care through some combination of four mechanisms, and each has a characteristic effect on a chronic, stigmatised, outpatient-dominant illness.

MECHANISM	HOW IT WORKS	EFFECT ON CHRONIC PSYCHIATRIC CARE
Tax-financed direct provision	Government funds and runs facilities; care is free or nominal at the point of use	Best fit in principle, because it can cover continuing outpatient care; limited in practice by staffing and reach
Publicly financed insurance	Government buys defined packages of care for defined beneficiaries from empanelled providers	Fits admissions and procedures; fits continuing outpatient treatment poorly unless deliberately designed to
Private voluntary insurance	Individuals or employers buy a product from an insurer, subject to regulation	Historically excluded or restricted psychiatric cover; the reason a parity requirement was needed at all
Out-of-pocket payment	The household pays at the point of use	Default when the other three fail; the mechanism through which illness produces impoverishment

Read the fourth row as the residual. Out-of-pocket payment is not a policy choice anybody makes; it is what happens when the first three do not cover something. Every argument for parity, for benefit-package design and for public financing in mental health is ultimately an argument about shrinking that fourth row.

RULE **An insurance benefit is defined by the event it pays for, and psychiatric care is mostly not an event.** Hospitalisation insurance pays for admissions. Chronic psychiatric care is a long sequence of inexpensive contacts with occasional expensive ones. A benefit designed around discrete admissions can be generous, well administered and still leave the ordinary course of a mental illness entirely unfunded.

9.2 The machinery of a publicly financed insurance scheme

To criticise a scheme intelligently you have to know where its joints are. A publicly financed hospitalisation scheme has six moving parts, and a mental health service can fail at any of them.

- **Beneficiary identification.** Somebody has to be established as entitled before care, which requires a document, a database entry and usually literacy or an intermediary. People who are homeless, wandering, undocumented or estranged from family fail here first, and that population overlaps heavily with severe mental illness.
- **Empanelment.** Only listed facilities can provide the benefit. If the psychiatric facilities in a district are not empanelled, the entitlement is real and unusable.
- **Benefit package design.** The scheme pays for named packages. What is not named is not paid for, and naming psychiatric care in package terms is genuinely hard because much of it is time rather than procedure.
- **Pre-authorisation.** Approval before treatment suits planned surgery and suits acute psychiatric admission badly, since the admission decision is often urgent and made out of

hours.

- **Package rates.** A fixed price per package determines whether providing the care is financially survivable for the facility. A rate set below cost produces empanelled facilities that quietly decline the cases.
- **Claims and audit.** Payment follows documentation. Psychiatric documentation is narrative rather than numerical, which makes claims slower to settle and easier to reject.

Notice that four of the six failure points are invisible to a patient. A family that is told the scheme does not cover this is rarely told which joint failed, and the distinction matters enormously for anybody trying to fix it.

9.3 Parity, and its three separable limbs

Parity is a word that hides three different demands, and an answer that treats it as one has lost most of the available marks.

- **Coverage parity.** Mental illness is included at all, rather than excluded by a clause. This is the crudest limb and the easiest to verify.
- **Benefit parity.** Where it is included, the limits, sub-limits, waiting periods, co-payments and caps are no less favourable than those applied to physical illness. A product can satisfy coverage parity and fail here completely, and this is where formal compliance most often stops.
- **Administrative parity.** The pathway to claiming is no harder: the same documentation burden, the same pre-authorisation friction, the same likelihood of settlement. This limb is almost never measured and is where a formally compliant product can still deliver very little.

The value of separating them is diagnostic. When you are told that a market has achieved parity, the follow-up questions are which limb, measured how, and by whom. A regulator can require the first with a stroke of a pen; the second requires product-level scrutiny; the third requires claims data that nobody publishes.

TRAP **Treating a regulatory requirement as evidence that products changed.** A direction to insurers is an input. Whether the products in the market now carry psychiatric cover on comparable terms, and whether people actually claim under it, are two further questions that require product analysis and claims data. Candidates routinely answer the second and third by asserting the first more loudly.

9.4 The regulator's instrument, which this chapter could not open

Here the honest position has to be stated in full, because this section's central document is one this chapter does not hold. The insurance regulator is the body that turns a statutory entitlement into a requirement on insurance products, and it does so through circulars addressed to insurers. The specific circular concerning mental illness coverage, its reference number, its date and its verbatim operative requirement, was sought for this chapter through the regulator's own health department circular index, through its master circular on health insurance business, and

through searches for the specific circular references that circulate in secondary discussion. It was not obtained.

What follows from that is a set of refusals, and they are deliberate. This chapter prints no circular number, no circular date and no quotation of what the regulator directed. It does not paraphrase the requirement, because the strength of a regulatory obligation lives in its modal verb and a paraphrase of an unread document invents that verb. It does not state when the requirement took effect. And it does not assert that any particular class of product must now behave in any particular way.

The corresponding examination answer is available and is a good one. A resident can say that the statutory right sits in the mental health legislation, that the regulator responsible for translating it into product requirements is the Insurance Regulatory and Development Authority of India, that this translation is effected through circulars to insurers, and that they would cite the specific circular by number and date rather than from memory. That answer is complete, correctly attributed and contains nothing invented.

9.5 Utilisation, and the rate that does not exist

Suppose the requirement is in force and products comply. The question that then matters for a population is how much psychiatric care is actually being funded through insurance. Answering it needs four things: a numerator of claims or of people treated, a denominator of eligible insured lives, a reference period, and the resulting rate.

No such published utilisation figure was obtained for this chapter. Searches of the regulator's circular index, its health department material and its master circular on health insurance business did not yield a numerator, a denominator or a rate for mental illness coverage. The consequence is that no proportion, no growth trend and no comparison with physical illness utilisation appears anywhere in this section.

The same applies on the public side. A publication of the national authority that administers the publicly financed scheme, specifying which mental health services are covered and, if published, their national utilisation, was sought and not obtained. So this chapter names the scheme and the machinery, and does not list its psychiatric packages or quote their uptake.

Sought THE REGULATOR'S CIRCULAR, ITS NUMBER, DATE AND WORDING	**Sought** A UTILISATION NUMERATOR, DENOMINATOR AND RATE	**Sought** THE COVERED-SERVICES LIST FOR THE PUBLIC SCHEME	**None held** SO NO FIGURE IS PRINTED IN THIS SECTION

Four tiles, three searches, no figures. That is an uncomfortable result for a section about financing and it is reported rather than concealed, because the alternative is worse. Financing figures are the easiest numbers in this whole chapter to half-remember from a seminar slide, and a coverage percentage attached to the wrong denominator is a fabrication that no digit check will catch.

9.6 What a resident should actually be able to say

The absence of figures does not make this material unanswerable; it changes what a good answer consists of. Structure it in four moves.

- **Name the layers.** A statutory right to parity, a regulator that translates it into product requirements, a publicly financed hospitalisation scheme for a defined beneficiary population, and out-of-pocket payment as the residual.
- **Say where mental illness fits badly and why.** Outpatient-dominant, longitudinal, time-based rather than procedure-based, and expensive in aggregate through frequency rather than through unit cost.
- **Separate the three limbs of parity** and say which of them any given claim about compliance actually addresses.
- **State the evidence position honestly.** That the operative regulatory instrument and the utilisation data were sought from the regulator and the national authority and are not available to you in citable form, and that you would not quote a coverage or utilisation figure you cannot attribute.

MNEMONIC

Identify, empanel, package, authorise, price, pay

The six joints of a publicly financed insurance scheme, in the order a case passes through them. Ask which joint failed before concluding that a scheme does not cover psychiatry.

PITFALL

Reassuring a family that a scheme will cover an admission before checking the two joints that most often fail: whether the facility is empanelled for this care, and whether the beneficiary record is in order today rather than in principle. Both fail silently and both fail at the point of admission, when the family has already committed. Check them at the outpatient visit, not at the door of the ward.

Parity is three separate demands, coverage, benefit and administration; a health system that has met the first has usually not been tested on the other two, and in India the instruments that would test them were sought for this chapter and are not available in citable form.

Epidemiology, burden and treatment gap

10 · National Mental Health Survey of India: prevalence figures and key findings

Every service argument in Indian psychiatry eventually rests on one survey, and most people who quote it cannot say what it counted. The **National Mental Health Survey** is the country's reference estimate of how much mental illness there is, who has it, and how much of it is going untreated. Its figures are the denominators behind almost every planning claim in this chapter. They are also the figures most often quoted at the wrong value, for the wrong population, without the year they refer to, and against a denominator that belongs to something else.

This section teaches the survey as a planning instrument: what it sampled, how it defined a case, what it found, and what those findings can and cannot be used to decide. How the estimates were

derived and what they mean epidemiologically are taught in the Research Methods, Biostatistics and Epidemiology notes. The treatment gap as a construct, its determinants and what closes it are taught in their own section of this chapter, and the burden of mental disorders expressed in disability-adjusted life years is taught in another; no figure from either is reproduced here.

10.1 What a national survey is for, and what it is not

A national survey produces a population estimate. A service produces a service statistic. The two answer different questions and are constantly confused, so fix the difference before anything else. A service statistic tells you about the people who arrived: how many were seen, what they had, what was done. It is silent about everyone who did not arrive, and in mental health that silent group is the larger one. A population survey goes out and asks, so it sees people who never present at all. This is why a national survey is the only instrument that can tell you the size of the problem rather than the size of the workload.

What it cannot do is tell you about a district. Survey estimates are constructed for the populations the design was powered to describe. Applying a national figure to a particular district, or a state figure to a taluka, imports an assumption about similarity that the survey never tested. This is the single commonest analytical error made with these numbers in Indian service planning.

RULE **A prevalence figure is a statement about a defined population in a defined period, measured with a defined instrument under a defined case definition.** Change any one of those four and the number changes. Quote a prevalence without them and you have quoted a rumour that happens to be accurate.

10.2 The design, and why the state list matters

The survey was undertaken during **2015-16**. Its fieldwork reached **39,532 individuals**, sampled across 720 clusters from 80 talukas in 43 districts of the 12 selected states. It covered 12 states, namely Punjab, Uttar Pradesh, Tamil Nadu, Kerala, Jharkhand, West Bengal, Rajasthan, Gujarat, Madhya Pradesh, Chhattisgarh, Assam and Manipur, drawn to represent the regions of the country.

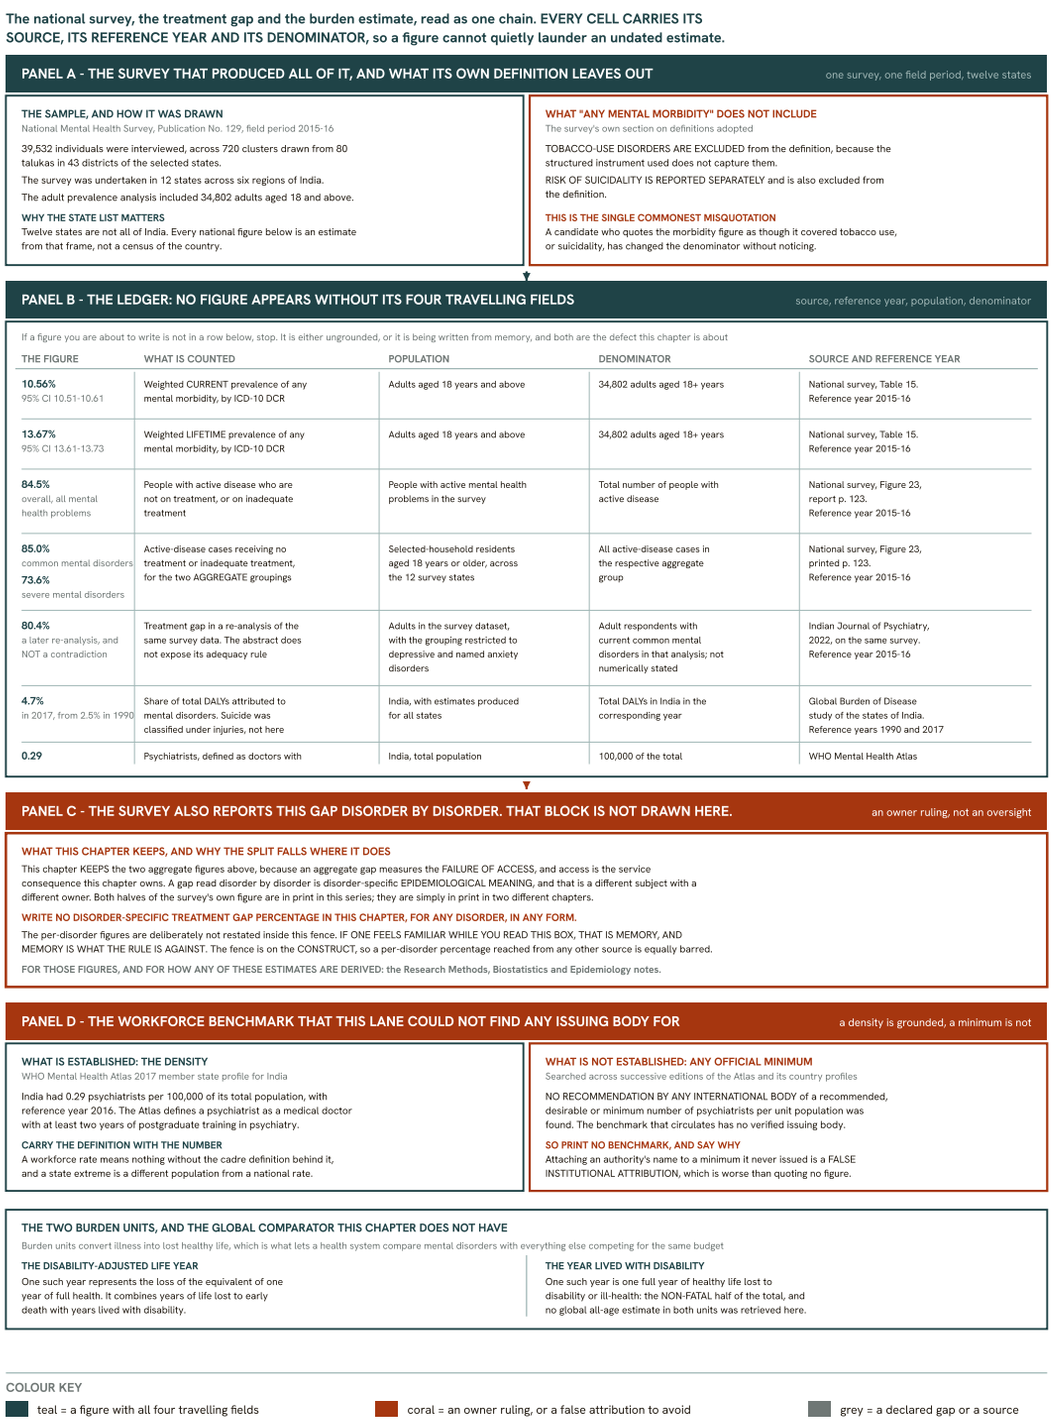


Fig 39.4 The survey, the gap and the burden, with every figure carrying its four travelling fields. The national survey interviewed 39,532 individuals across 720 clusters in 43 districts of 12 states, and its morbidity definition excludes tobacco-use disorders and reports risk of suicidality separately. Current prevalence was 10.56 per cent and lifetime 13.67 per cent. The treatment gap is printed here only in its AGGREGATE forms: 84.5 per cent overall, 85.0 per cent for common and 73.6 per cent for severe mental disorders, with 80.4 per cent from a later re-analysis whose disorder grouping was narrower. The survey's disorder-by-disorder block is owned by the Research Methods, Biostatistics and Epidemiology notes and no per-disorder percentage appears on this plate. Mental disorders accounted for 4.7 per cent of total DALYs in India in 2017, against 2.5 per cent in 1990. India had 0.29 psychiatrists per 100,000 population, and no international body's recommended minimum could be found, so none is printed.

The accompanying figure sets out the chain these findings sit in, from the survey's prevalence estimates through the treatment gaps derived from them to the burden expressed as health loss, with the source, reference year and denominator printed against every cell so that no figure in the chain can be read without them.

Read the state list rather than skipping it, because it is the survey's most important limitation and its most quotable one. Twelve states were sampled out of a much larger number, deliberately spread across the regions of the country so that the whole could be described from the parts. That is a legitimate and standard design. It also means that a state not on that list has no estimate of its own from this survey, and that any national figure derived from it is a construction from twelve sampled populations rather than a measurement of every state.

TRAP **Quoting the reference year as the year you are reading it.** These are figures from fieldwork conducted in 2015-16, and the reference year travels with them permanently. An estimate does not become current by being repeated, by being reprinted in a later document, or by being cited in a recent policy paper. Every time you quote one of these numbers, say the period it describes.

10.3 The case definition, which is where most misquotation begins

If you learn one thing from this section that other candidates will not know, make it this. The survey's headline category, any mental morbidity, is a constructed definition with explicit inclusions and, more importantly, explicit exclusions. Diagnoses were made against ICD-10 criteria and captured through a structured diagnostic interview, the Mini International Neuropsychiatric Interview, which fixes both what could be detected and what could not.

Two exclusions in particular change what the headline number means. **Tobacco-use disorders were excluded from the definition, and risk of suicidality, although captured, was reported separately and also excluded from it**. Both decisions are defensible and both are consequential. Tobacco dependence is extremely common, so including it would have raised the headline figure substantially and would have made it a different sort of number. Risk of suicidality is a risk state rather than a disorder, so folding it into a morbidity total would have mixed two kinds of thing.

The consequence for a reader is precise. When somebody quotes the national prevalence of mental disorders in India, the correct follow-up question is which construct they mean, because this headline excludes the commonest substance-use problem in the country and excludes the risk state that most service planning around suicide is directed at.

IN INDIA

Do not size a tobacco cessation component or a suicide prevention component from this survey's headline morbidity figure. Tobacco-use disorders are excluded from the definition and risk of suicidality is reported separately and excluded from it, so the headline number contains neither. A district plan that budgets for either on the basis of the overall morbidity estimate has sized it from a denominator that deliberately leaves it out, and will under-provide in exactly the two areas where demand is highest.

10.4 The headline findings, with their denominator attached

The adult prevalence analysis was conducted on **34,802 adults** aged 18 years and above. That figure is the denominator of the prevalence estimates, and it is smaller than the total number of people interviewed because the adult analysis is a subset of the whole sample. Carry it with the percentages; a proportion without its denominator cannot be interrogated.

Among adults aged 18 years and above, the weighted current prevalence of any mental morbidity was 10.56 per cent, with a 95 per cent confidence interval of 10.51 to 10.61. The corresponding weighted lifetime prevalence of any mental morbidity was 13.67 per cent, with a 95 per cent confidence interval of 13.61 to 13.73. Both are weighted percentages, which is stated in the report's own table footnote and is not a detail: weighting is what allows a sample drawn unevenly across strata to speak for the population it was drawn from.

The survey also reports how much of that morbidity was going untreated. Its treatment gap for mental health problems was **84.5 per cent**. Read the denominator carefully, because this is where the number is most often mangled: it is a proportion of the total number of people with active disease, counting people with active disease who are not on treatment or on inadequate treatment. It is not a proportion of the population. The treatment gap as a construct, its determinants and the interventions that narrow it are taught in their own section of this chapter, and the gap read disorder by disorder, with its epidemiological meaning, is taught in the Research Methods, Biostatistics and Epidemiology notes.

39,532	10.56%	13.67%	84.5%
INDIVIDUALS INTERVIEWED ACROSS THE SELECTED STATES	CURRENT ANY MENTAL MORBIDITY, ADULTS AGED EIGHTEEN AND OVER	LIFETIME ANY MENTAL MORBIDITY, SAME ADULT SAMPLE	TREATMENT GAP, OF PEOPLE WITH ACTIVE DISEASE

10.5 Four numbers, four denominators

Set the headline figures side by side with what each is a proportion of, and the difference between them stops being memorisable trivia and becomes a working tool.

FIGURE	POPULATION IT DESCRIBES	WHAT IT IS A PROPORTION OF	WHAT IT DOES NOT TELL YOU
Current prevalence 10.56 per cent	Adults aged 18 years and above in the sampled states	34,802 adults in the analysis	How many will develop illness later, or how many are in treatment
Lifetime prevalence 13.67 per cent	The same adult sample	34,802 adults in the analysis	How many are unwell now, which is the number a service needs
Treatment gap 84.5 per cent	People with active mental health problems in the survey	The total number of people with active disease	Why they are untreated, which is a separate section of this chapter
39,532 individuals interviewed	The whole sampled population, adults and others	Not a proportion; a sample size	Anything about states outside the sampling frame

The second row repays attention. Lifetime prevalence is always the larger and more quotable figure, and it is the less useful of the two for planning a service, because a service treats people who are unwell now. Where a planning document quotes a lifetime figure to justify current establishment, it has quietly inflated its case.

10.6 Children, and the limit this survey states about itself

The adult findings above are exactly that: adult findings, from an adult analysis, with an adult denominator. The survey carried an adolescent component, but it was a feasibility sub-sample confined to some of the survey states rather than a nationally representative adolescent survey, and it does not support an all-India childhood estimate.

The adolescent figure, the limits attached to it and the position that no established all-India psychiatric prevalence can be quoted for younger children are taught in the Child psychiatry: assessment, treatment, services and protection notes. No child or adolescent prevalence figure appears in this chapter, and no adult figure here is extended downwards to children. Extrapolating an adult rate to a child population is not a conservative approximation; it is a different claim about a different population made without evidence.

10.7 Using the findings for priority-setting without misusing them

The reason to know the design so well is that priority-setting arguments are won and lost on the denominator. Four disciplines make the difference between a survey-based argument that holds and one that collapses under a single question.

- **Name the construct before the number.** Any mental morbidity, current, adults, with tobacco-use disorders excluded and suicidality reported separately, is a specific thing. Saying "the prevalence of mental illness in India" and then giving a figure omits every part that makes it true.
- **State the reference period every time.** The estimates describe fieldwork in a stated period and nothing since. If you are asked how much of this is still true, the honest answer is that the survey cannot say, because a survey speaks only about when it was done.
- **Keep current and lifetime apart.** Use the current figure for establishment and workload arguments and the lifetime figure for arguments about population exposure over a life course.
- **Never mix denominators.** A prevalence is a proportion of a population, a treatment gap is a proportion of the people with the condition, and no arithmetic combining the two directly produces anything meaningful.

MNEMONIC

Who, when, how defined, out of what

The four fields that make a survey number a fact. Which population, which reference period, which case definition and instrument, and which denominator. Any of the four missing, and the figure cannot be checked by anybody who hears it.

PITFALL

Presenting the national estimate to a district administration as though it described their district. The survey sampled a defined set of states through a clustered design, and the national figure is built from those samples. A district that differs in its age structure, its urban proportion, its migration pattern or its substance-use profile will differ in its true prevalence, and nobody knows in which direction. Use the national estimate to argue that a service is needed, and local data, however imperfect, to argue how much.

Current weighted prevalence of any mental morbidity in adults was 10.56 per cent and lifetime was 13.67 per cent, both from fieldwork in 2015-16, on a definition that excludes tobacco-use disorders and reports suicidality separately.

11 · Treatment gap for mental disorders in India and its determinants

The treatment gap is the single most quoted number in Indian community psychiatry and the one most often quoted against the wrong denominator. It is not a measure of how much illness there is. It is a measure of how completely a health system is failing to reach the people who already have it. That distinction is the whole of this section: everything that follows is about what makes the gap the size it is, who is left on the wrong side of it, and what has actually been shown to move it.

How the estimates are derived, and what the gap means read disorder by disorder, are taught in the Research Methods, Biostatistics and Epidemiology notes. The national survey that produced the Indian estimates, its design and its own overall reported figure are taught in the section of this chapter that owns the survey. What belongs here is the causal and service account: determinants, inequities, consequences and interventions.

11.1 What the gap measures, and out of what

A **treatment gap** is the proportion of people with a condition who are not receiving treatment, or who are receiving treatment inadequate to their condition. Three features of that definition do all the work and each is a route to error.

- **The denominator is the diseased, not the population.** A gap of four fifths means four of every five people who have the condition are untreated. It says nothing about how common the condition is. Multiplying a gap by a population is arithmetic on two unrelated things.
- **Inadequate treatment counts as gap.** The measure is not simply contact. A person prescribed a subtherapeutic dose, seen once and never followed up, or treated by someone unqualified to treat them, sits inside the gap despite having reached a service. This is why closing a gap is harder than opening clinics.
- **The gap is population-specific.** It differs by disorder, by severity, by state, by sex and by wealth, so a single national number is an average across all of those and describes nobody exactly.

RULE **Never quote a treatment gap without naming the population it belongs to.** The commonest defect in this whole area is a figure carried across from the population it was calculated for to a different one. A correctly transcribed percentage attached to the wrong group is still a false statement, and it is one that no arithmetic check will catch.

11.2 The national figures, and the values that do not agree

India's national survey reports the gap for two aggregate groupings. **The treatment gap was 85.0 per cent for common mental disorders and 73.6 per cent for severe mental disorders**. The population is selected-household residents aged 18 years or older across the 12 survey states, the reference period is 2015-16, and the denominator in each case is all active-disease cases in the respective common-mental-disorder or severe-mental-disorder group. What is being counted is active-disease cases receiving no treatment or inadequate treatment.

Notice what the pair tells you before you go any further. The gap is *larger* for the common disorders than for the severe ones. That inversion surprises people who expect severity to track access, and it is one of the most instructive facts in this chapter. Severe illness is disruptive, visible to families and hard to ignore, so it eventually generates a contact even where services are poor. Depression and anxiety are survivable in a way that lets a person and a household absorb them indefinitely without anyone deciding that this is illness.

A lower figure circulates widely for the same construct, and a resident who meets it should be able to place it. A later analysis based on the same national survey data reported a treatment gap of **80.4 per cent** for its own analysis of common mental disorders. It is not a competing measurement of the same thing: that analysis restricted common mental disorders to depressive disorders and six named anxiety disorders, so its denominator is adult respondents with current common mental disorders included in that analysis and not the survey's own broader grouping. The opened abstract reports the treatment gap but does not expose its detailed numerator, its treatment-adequacy rule or its treatment window, and does not state the numeric denominator.

A third figure for common mental disorders also circulates and is not printed here. It was sought directly, in the published analysis and through several routes to its own tables, and it could not be established whether it is an overall pooled estimate or a subgroup estimate. A percentage whose population cannot be determined is not a fact yet, and this chapter does not print it in either reading.

TRAP **Treating two treatment gap figures from one survey as a contradiction to be resolved.** They usually differ because the analyses defined the disorder grouping differently, applied different adequacy rules, or used different time windows. The right response is not to pick one and dismiss the other; it is to state which grouping each describes. A candidate who says "the survey's own figure for common mental disorders, against a later re-analysis restricted to depressive and named anxiety disorders" has demonstrated exactly the discipline the examiner is testing.

11.3 What is deliberately not printed here

The same survey also reports the treatment gap disorder by disorder. Those per-disorder percentages are not printed anywhere in this chapter. They are neither missing nor doubtful: they

are grounded, verified against the survey's own report, and taught, with their derivation and their epidemiological meaning, in the Research Methods, Biostatistics and Epidemiology notes. A per-disorder gap figure is a statement about the epidemiology of that disorder, which is that chapter's subject; an aggregate gap is a statement about access failure, which is this one's.

That boundary is worth stating explicitly rather than silently observing, for a practical reason. If a per-disorder figure feels familiar as you read this, that familiarity is memory rather than evidence, and memory is exactly what the rule exists against. Write the aggregate figures, say that the survey reports the breakdown, and send the reader to the chapter that carries it.

11.4 Demand-side determinants, and the first person a family actually consults

A gap that large is not produced by one barrier; it is produced by a sequence, and the sequence begins before anybody has thought of a clinic. The national survey itself identifies the demand-side conditions behind it: **limited awareness, sociocultural beliefs, values and stigma** are named as demand-side barriers contributing to the treatment gap.

Qualitative work in Indian states fills that in from the point of view of the people making the decision. In a situational analysis of the mental health pathway in one southern state, stakeholders including policymakers, health professionals, non-governmental organisation staff and clients identified stigma, knowledge and expectations, lack of practical treatment access, and cost as influences on decisions about seeking care, drawn from 10 semi-structured interviews and 5 focus groups, each focus group containing 6 to 8 participants.

Stigma is not merely one item on that list; in some populations it is the decisive one. Among 100 consecutive included patients in a Kerala study of people who had attended a faith healer before reaching tertiary psychiatric care, **44 per cent** reported stigma of psychiatric illness as a main reason for avoiding psychiatric treatment.

The pathway data show where people go instead. In a South Indian tertiary outpatient sample of 104 patients and informants, the first carer consulted was an allopathic practitioner for 51.0 per cent, a religious healer for 26.9 per cent, and a mental-health professional for 22.1 per cent. Read those three proportions as a service design brief rather than as a curiosity. Roughly half of first contacts were with a doctor who was not a psychiatrist, which is the strongest available argument for integrating mental health into general and primary care. Around a quarter were with a religious healer, which is the argument for the partnership work taught in this chapter's section on traditional and faith-based care. Fewer than a quarter reached a mental health professional first.

IN INDIA

Design the first contact around where people actually go, not around where you wish they went. In the pathway sample above, 51.0 per cent of first carers were allopathic practitioners and 26.9 per cent were religious healers, against 22.1 per cent mental-health professionals. The service decision that follows is to invest in the recognition and referral capacity of general practitioners and in structured, respectful working relationships with faith-based providers, because both of those sit upstream of any psychiatric service and will continue to whatever the psychiatric service does. Note the limit as you use it: this is one tertiary outpatient sample of 104 patients and informants, not a national estimate.

11.5 The supply side, and the benchmark that does not exist

Demand-side barriers do not act alone. Situational work in northern India identified a dearth of mental-health professionals and inadequate services and facilities as adversely affecting the mental health treatment gap, drawn from 13 purposively recruited interview participants, plus local health-facility records and policy documents.

The workforce position can be put more precisely. For India, **0.29 psychiatrists per 100,000 total population** is reported in WHO Mental Health Atlas 2017, for the reference year 2016, against a denominator of 100,000 of India's total population using the official population estimate shown in the country profile. The definition matters as much as the value: the figure counts a medical doctor who has had at least two years of postgraduate training in psychiatry, which is the Atlas's own workforce definition and not necessarily the definition another source would use.

Now the part that most teaching gets wrong. It is conventional to follow a density figure with a recommended minimum, usually attributed to an international body. No such recommendation could be established for this chapter. An explicit recommended, desirable or minimum number of psychiatrists per lakh of population attributable to the World Health Organization or to the refugee agency was sought across the Atlas report and country profiles for several editions, the global health observatory indicator metadata, the intervention guide, and both agencies' own pages, and it was not found.

The consequence is a refusal that a candidate can make confidently: this chapter prints no workforce benchmark. Attributing a minimum ratio to an international body that has not issued one is a false institutional attribution, and it is worse than printing no benchmark at all, because it borrows an authority that was never given. The defensible sentence is that the density is low in absolute terms, that it is reported by a named body for a stated reference year against a stated denominator, and that no internationally issued minimum could be verified to compare it against.

85.0%	73.6%	0.29	None found
TREATMENT GAP, COMMON MENTAL DISORDERS, OF ACTIVE CASES	TREATMENT GAP, SEVERE MENTAL DISORDERS, OF ACTIVE CASES	PSYCHIATRISTS PER LAKH OF TOTAL POPULATION, ATLAS REFERENCE YEAR 2016	VERIFIED INTERNATIONAL MINIMUM WORKFORCE RATIO

11.6 Distance against money: what actually predicts reaching care

The intuitive model of access is geographical, and the best Indian evidence available here does not support it. In a rural implementation area in Madhya Pradesh, road distance to the nearest public provider of depression treatment was not associated with the probability of seeking treatment: the odds ratio was 1.00, with a 95 per cent confidence interval of 0.98 to 1.02, among 568 adults with probable depression.

In a population-based survey in the same district, what did predict contact coverage was economic position. Among 576 adults with probable moderate to severe depression within a 3,220-adult baseline survey, higher-income occupation was associated with a prevalence ratio of 7.17, with a 95 per cent confidence interval of 1.23 to 41.70, and land ownership with a prevalence ratio of 1.83, with a 95 per cent confidence interval of 1.21 to 2.76 for higher contact coverage.

Read those two findings together and the model of access changes. Building a clinic nearer to people does not, by itself, appear to be the binding constraint; being poorer and owning nothing appears to be. Treat the wide interval on the occupation ratio with the caution it deserves, and note that both studies concern one district and depression rather than the country and mental illness in general. Even with those limits, the direction is a corrective to the standard assumption, and it points service design at cost, lost wages and social position rather than at distance alone.

What is not available is the national quantification of that inequity. Quantified Indian treatment-gap or contact-coverage inequities comparing rural with urban populations, or comparing groups by sex or by education, without resorting to the disorder-specific figures owned by another chapter, were sought in the national survey report, in a summary of it, and in each of the primary Indian studies named above. None was obtained. So the inequity is taught here as a mechanism supported by district-level studies, and no national inequity percentage is printed.

11.7 What closes the gap, and what the trials actually measured

India has produced an unusually strong body of trial evidence on delivering psychological treatment through non-specialists. It is worth knowing well, and it is worth knowing precisely, because almost every trial measured a symptom outcome rather than a coverage outcome, and the distinction is the difference between showing that a treatment works and showing that a gap has closed.

STUDY AND SETTING	RESULT	WHAT WAS MEASURED
Lay health worker led collaborative care in primary care, Goa	A 30 per cent reduction between arms in the prevalence of an ICD-10 common mental disorder diagnosis	A symptom-diagnosis outcome across follow-up, not contact coverage
Brief psychological treatment for depression added to enhanced usual care, Goa	Remission 63 per cent against 48 per cent	Participants meeting the trial's remission criterion; a symptom outcome, not treatment contact or population coverage

STUDY AND SETTING	RESULT	WHAT WAS MEASURED
Brief counselling for harmful drinking added to enhanced usual care, Goa	Remission 54.3 per cent against 31.9 per cent, adjusted prevalence ratio 1.71	A trial-defined remission outcome, not contact coverage
Community champion led psychosocial intervention, Gujarat	Recovery odds ratio 2.2, with a 95 per cent confidence interval of 1.2 to 4.6	Categorical recovery from depression or anxiety symptoms; not service contact or population coverage
Peer-delivered perinatal depression intervention, Goa	Six-month remission 73 per cent against 60 per cent	A trial remission definition assessed six months after childbirth; not contact coverage
Anti-stigma and mental health service delivery in rural primary care, Andhra Pradesh	94.5 per cent of the intervention group against 2.5 per cent of controls self-reported receiving mental-health treatment at 12 months	Self-reported receipt of treatment, not independently verified treatment, symptom remission or whole-population coverage
Grass-roots community programme, rural Maharashtra	Associated with a substantial increase in equitable contact coverage for depression and improved mental health literacy	Contact coverage and literacy indicators before and after implementation

Work down the third column, because that column is the point. Five of the seven measured whether people got better, in populations already identified and enrolled. Only the last two measured whether people were reached at all, and one of those is described qualitatively here because its exact baseline and follow-up coverage figures could not be retrieved from any accessible copy of the original report. So the Indian evidence that non-specialist delivered treatment works is strong. The Indian evidence that any of it has closed a population gap is much thinner, and rests largely on those two.

The control arm figure in the rural Andhra Pradesh trial deserves separate attention. 2.5 per cent of controls self-reported receiving mental health treatment over twelve months, in a high-risk cohort identified by screening. That is the treatment gap seen from inside a trial, in people who had already been identified as needing care, and it is the most concrete illustration in this section of how little happens by default.

TRAP **Citing a remission rate as evidence that an intervention closes the treatment gap.** Remission is measured in people who were enrolled, and enrolment is the very step the gap consists of failing at. An intervention can produce excellent remission in everyone it reaches and change population coverage not at all. When you are asked what closes the gap, name the outcome the study measured before you name the effect size.

11.8 The evaluation nobody has published, and how to answer this

The standard answer to the treatment gap in India is the district programme. An official Indian evaluation reporting a measured before-and-after or controlled change in treatment gap attributable to that programme was sought for this chapter from the ministry, the directorate, the national health mission and the official assessment of the mental health system. None was

obtained: programme objectives, approved districts, administrative activity and service reporting are all documented, and an attributable effect on the gap is not.

That absence is the most useful single thing in this section. The country's principal instrument for closing the gap has not been shown, in any citable official evaluation, to have changed the gap. That is not a claim that it has failed; it is a claim that the question has not been answered, and the reason is that programme reporting was designed to record activity rather than to support attribution.

MNEMONIC

Know, seek, reach, receive, continue

The five steps between having a mental illness and being adequately treated. The gap is the sum of everyone lost at each step, which is why an intervention aimed at one step alone rarely moves the total.

PITFALL

Assuming the gap will fall as services expand. Expansion adds capacity at the reach step and does nothing for the know, seek or continue steps, which is why services can grow while the gap stays where it is. The corollary is practical: if you are given resources to reduce a gap, spend some of them on identification and on retention in care, not all of them on capacity, because capacity that nobody enters is invisible in every outcome measure that matters.

The treatment gap is a proportion of people with the condition, not of the population: 85.0 per cent for common mental disorders and 73.6 per cent for severe mental disorders in India's national survey, and the higher figure belongs to the milder group.

12 · Global and Indian burden of mental disorders: DALYs and YLDs

Counting people with an illness tells a finance ministry almost nothing, because it makes a fatal cancer and a mild phobia the same size. Health accounting solves that by measuring what is lost rather than who is affected, and the instrument it uses is the disability-adjusted life year. This is the measure through which psychiatry competes for resources against cardiology, obstetrics and infectious disease, and a resident who cannot explain how it is built cannot make the case for a mental health service in the room where the money is allocated.

This section teaches the metric and the Indian result. The definitions of incidence, prevalence and the measures of association, and the study designs behind them, are taught in the Research Methods, Biostatistics and Epidemiology notes. The national survey's prevalence findings and the treatment gap have their own sections in this chapter and no figure from either is reproduced here.

12.1 From counting cases to counting lost health

A case count answers how many. It cannot answer how bad, how long, or how much of a life is affected, and those are the questions any allocation decision turns on. Burden metrics were built to make different diseases commensurable by converting everything into a single currency: time lived in less than full health, plus time not lived at all.

The disability-adjusted life year is that currency. **One disability-adjusted life year corresponds to one year of full health lost, and the measure combines years of life lost through premature mortality with years lived with disability**. It therefore adds a severity dimension to a prevalence count, which is exactly what a prevalence count lacks.

Its two components answer different questions and they behave differently.

- **Years of life lost** capture mortality: how many years of expected life a death removed. A disease that kills young generates a great deal of this.
- **Years lived with disability** capture morbidity: how long people live with the condition, weighted by how disabling it is. **One year lived with disability represents one full year of healthy life lost due to disability or ill health, with prevalent conditions weighted by their disability weights**.

RULE **A burden estimate is the product of how many, how disabling, and for how long.** Change any one factor and the total changes without the number of cases moving at all. This is why a condition can become a larger share of national burden while becoming no more common.

12.2 Why psychiatry lives almost entirely in one half of the measure

Mental disorders are, in health accounting terms, a morbidity problem rather than a mortality problem, and that single fact explains most of psychiatry's position in these tables. They begin early, often in adolescence or early adult life. They persist or recur for decades. They impair function substantially while people remain alive. So they accumulate years lived with disability at a high rate and contribute comparatively little through years of life lost.

Three consequences follow, and they are the examinable content of this subtopic.

- **Psychiatry looks small in mortality-based accounts and large in disability-based ones.** Any argument built on death rates alone will understate mental illness, which is why mortality-led health planning has historically neglected it.
- **Duration does the work.** A condition that begins at twenty and persists accumulates burden for decades. This is the quantitative form of the argument for early intervention: preventing or shortening a long illness saves far more health than treating a short one.
- **Disability weights are a modelling choice, not a measurement.** How disabling a condition is judged to be is derived through survey and modelling methods, and a change in weights changes every downstream total without anybody's health having changed. Treat a burden figure as a model output, not as a count.

12.3 The Indian result, and the classification that sits underneath it

The Indian state-level burden work reports the contribution of mental disorders to total health loss in the country and how it has moved. **Mental disorders accounted for 4.7 per cent of total disability-adjusted life years in India in 2017, compared with 2.5 per cent in 1990.** The published estimates carry uncertainty: 2.5 per cent with a 95 per cent uncertainty interval of 2.0 to 3.1 in 1990, and 4.7 per cent with an interval of 3.7 to 5.6 in 2017. The denominator in each case is total disability-adjusted life years in India in the corresponding year, and the estimates were produced for all states of India as part of the global burden of disease work for that cycle.

Now the qualification that changes how the figure should be read, and which most people who quote it have never seen. What is being counted is disability-adjusted life years attributed to mental disorders as defined in that burden cycle, in which suicide was classified under injuries. Self-harm and suicide therefore sit outside this share. Whatever burden India's suicide mortality represents, it is not inside the figure above.

That is not a criticism of the estimate; it is how the classification works, and the same convention is applied across countries so that the comparisons hold. It is, however, decisive for how you use the number. A mental health service argument built on this share is built on non-fatal psychiatric morbidity alone, and a candidate who states that limitation while quoting the figure has demonstrated the four-field discipline more convincingly than one who simply recites the value.

2.5%	4.7%	Not held
MENTAL DISORDERS AS A SHARE OF TOTAL INDIAN DALYS, 1990	THE SAME SHARE IN 2017, SUICIDE CLASSIFIED SEPARATELY	GLOBAL ALL-AGE ESTIMATE IN BOTH DALYS AND YLDS

IN INDIA

When you are making the case for a mental health service to an administrator, use the share-of-burden framing rather than a case count, because it is the currency the allocation decision is already made in: 4.7 per cent of total Indian disability-adjusted life years in 2017, against 2.5 per cent in 1990. Then say, in the same breath, that suicide was classified under injuries in that cycle and sits outside the figure. Volunteering the limitation is what makes the rest of your numbers credible in a room where somebody else is arguing for cardiology.

12.4 How to read a rising share without over-reading it

A share that nearly doubles looks like an epidemic and usually is not one. The share of total burden attributable to a condition can rise for at least five reasons, and only one of them is that more people are becoming ill.

- **Other causes fell.** A share is a ratio. When communicable disease, maternal and child mortality and nutritional deficiency burden decline, everything else rises as a proportion without changing in absolute terms. This is the dominant mechanism in most middle-income countries.

- **The population aged and grew.** More adults living longer means more person-years at risk of conditions that persist.
- **Detection and reporting improved.** Better data produce higher estimates for conditions that were previously invisible.
- **Methods changed.** Successive burden cycles revise disability weights, case definitions and models, and a revision can move a share on its own.
- **Incidence genuinely rose.** Possible, and the hardest of the five to demonstrate.

The disciplined sentence is therefore that the share of India's health loss attributable to mental disorders roughly doubled over that period, that this is consistent with an epidemiological transition in which other causes fell faster, and that a share cannot by itself establish that the underlying rate of illness rose.

TRAP **Reading a rising share of burden as a rising incidence of illness.** A proportion moves when either its numerator or its denominator moves. In a country where the denominator has been falling for other reasons, the numerator can be flat and the share can still climb. Say "share of total burden" every time, and the error becomes impossible to make.

12.5 Three ways to express the same burden

One population, one set of illnesses, three ways to describe them. Each is right for a different audience and each hides something.

EXPRESSION	WHAT IT ANSWERS	WHAT IT HIDES	WHO USES IT
Case count or prevalence	How many people are affected	Severity, duration and mortality; a mild and a severe case count the same	Service planners sizing establishment and workload
Years lived with disability	How much non-fatal health is being lost	Deaths entirely; a condition that kills quickly scores low	Advocates for chronic and disabling conditions
Disability-adjusted life years	Total health lost, fatal and non-fatal together	Who exactly is affected, and the modelling assumptions behind the weights	Finance and planning bodies comparing across disease areas

A complete argument uses all three in sequence: how many people, how much health is being lost, and therefore what share of the allocation is defensible. An argument that uses only the first is easy to dismiss on the ground that the conditions are not serious, and an argument that uses only the third is easy to dismiss on the ground that it is a modelling artefact.

12.6 The global figure this chapter does not print

This item is titled for the global and Indian burden, and no global figure appears in it. That is a declared evidence position rather than an oversight. A global, all-age, both-sex estimate for all mental disorders expressed in both disability-adjusted life years and years lived with disability, with a stated reference year and the corresponding global totals as denominators, was sought

from the international health agency and from the institute that produces the burden estimates, through their observatories and through targeted searches of both. It was not obtained in a form carrying all four fields.

Two further gaps sit alongside it and both concern access rather than existence. A serving address for the published Indian state-level burden analysis could not be reached: the hosted copy returned an error, a repair attempt returned the same dead address, and neither a keyword search nor an exact-title search of the biomedical databases resolved the article. Two separate records in this chapter's evidence ledger point at that same unreachable source. The Indian figures above are printed because they are carried with their full provenance in this chapter's own evidence, not because the article could be opened afresh.

The consequence for a candidate is a specific and defensible position: quote the Indian share with its reference years, its uncertainty intervals, its denominator and its suicide classification; and, if asked for the global figure, say that a global estimate carrying both metrics with their denominators was sought and not obtained in a citable form, rather than reaching for a percentage that is widely repeated and rarely sourced.

MNEMONIC

How many, how bad, how long, how fatal

The four inputs to any burden estimate. Prevalence supplies the first, disability weights the second, duration the third, and mortality the fourth. Psychiatry scores heavily on the first three and lightly on the fourth, which is precisely why it needs this metric to be visible at all.

PITFALL

Mixing burden figures from different estimation cycles in one argument. Successive cycles revise weights, definitions and models, so a figure from one release and a figure from another are not two observations of the same quantity, and the difference between them can be entirely methodological. Take every figure in a single argument from one declared release, and say which one.

Mental disorders accounted for 4.7 per cent of total Indian disability-adjusted life years in 2017 against 2.5 per cent in 1990, a share of total health loss in which suicide was classified under injuries and therefore does not appear.

13 · Basic epidemiological concepts in psychiatry (incidence, prevalence, study designs)

This chapter uses epidemiological measures on almost every page and defines none of them. Incidence, prevalence, the relation between them, relative risk, the odds ratio, attributable risk and the observational designs that produce all of these are taught in the Research Methods, Biostatistics and Epidemiology notes. That is where the definitions, the derivations and the worked arithmetic live, and they are not repeated here.

The division is worth stating precisely because it is unusually clean. Defining a measure is a methods question. Using a measure to decide what a health system should do is a public health question. This chapter is entirely the second, and the three sections immediately before this one are where the using happens: the national survey and its prevalence findings, the treatment gap and its determinants, and the burden of mental disorders expressed as disability-adjusted life years and years lived with disability.

13.1 Where each measure is actually applied in this chapter

A reader who wants to see these concepts working rather than defined should read the three sections named above in that order, because they form a chain and each link uses a different measure for a different decision.

- **Prevalence answers how many people have the condition now.** That is the number a district uses to forecast caseload and to size an establishment. It is applied in the section on the national mental health survey.
- **The treatment gap answers how many of those people are not receiving adequate care.** That is a proportion of the diseased population, not of the general population, and it measures the failure of access rather than the frequency of illness. It is applied in the section on the Indian treatment gap.
- **Burden answers how much health is being lost, weighted by severity and duration.** That is the number a finance ministry compares across diseases when it allocates. It is applied in the section on disability-adjusted life years and years lived with disability.

Read as a chain, those three answer three different questions that are constantly confused with one another in examinations: how common, how untreated, and how costly. They can move in opposite directions. A condition can be common and cheap, or rare and devastating, or common and almost entirely untreated, and each of those combinations implies a different service response.

RULE **Every measure in this chapter is quoted with what it is a proportion of, and that denominator is part of the number.** A prevalence is a proportion of a population. A treatment gap is a proportion of the people with the condition. A burden share is a proportion of total health loss. Three proportions, three denominators, and the commonest examination error in this whole area is carrying one number across to the other's denominator.

13.2 The boundary, and why holding it protects the reader

There is a specific reason this section refuses to give a short definition of prevalence and move on. The measures are written in the same words on both sides of the boundary, so a short restatement here would look identical to correct teaching and would be impossible to police mechanically. What makes it a defect is not the wording but the consequence: two definitions of one measure in one book eventually diverge, and a reader who has met both cannot tell which is authoritative. One owner, one definition.

TRAP **Treating a treatment gap as though it were a prevalence.** They are reported in the same units and in adjacent sentences of the same survey report, and their denominators are entirely different. A gap of the order of four fifths does not mean four fifths of the population is ill; it means that of the people who are ill, roughly four in five are not adequately treated. Say the denominator out loud before you quote the figure.

How common	How untreated	How costly
PREVALENCE, APPLIED TO CASELOAD FORECASTING	TREATMENT GAP, APPLIED TO ACCESS	BURDEN, APPLIED TO ALLOCATION

MNEMONIC

Common, untreated, costly

Three questions, three measures, three denominators, three different service decisions. The measures themselves are defined in the Research Methods, Biostatistics and Epidemiology notes; this chapter only ever uses them.

Incidence, prevalence, measures of association and the observational designs that yield them are taught in the Research Methods, Biostatistics and Epidemiology notes; this chapter applies them and defines none of them.

Suicide prevention and farmer's suicide

14 · Farmer's suicide in India: epidemiology, determinants and prevention

Farmer suicide is the one public health statistic in India that everybody has heard and almost nobody can state correctly. It is quoted in politics, in journalism and in examinations, usually as a bare count and usually without the two things that would make it interpretable: what the counting category actually contains, and what the figure is a proportion of. This section teaches the number properly, and it teaches why the most important quantity in this whole area, the rate at which people in farming die by suicide compared with everybody else, cannot currently be computed at all.

The farming sector as a recognised at-risk group within general suicide epidemiology, and the population prevention levers that apply to it, are taught in the Somatic-symptom, sexual, impulse-control disorders and suicide notes. The national suicide prevention strategy, means restriction as a programme, and the epidemiology of suicide in India generally have their own sections in this chapter. What belongs here is the occupational count itself, its classification, its denominator problem and what would count as evidence that anything has worked.

14.1 Why an occupational category is a public health object

Suicide surveillance in India runs through a police-recorded route. The National Crime Records Bureau (NCRB) compiles cases reported through the police system and publishes them classified in several ways, one of which is by profession. That classification is what turns a distributed

personal tragedy into a countable population, and it is the reason a farming category exists in the national data at all.

An occupational category is worth having for three reasons. It identifies a population that shares an exposure, which is what makes targeted prevention possible. It allows a trend to be watched over time. And it creates political visibility, which in this case has been considerable. It also carries three costs, and the whole of the rest of this section is about them: the category has to be operationally defined, the people in it have to be correctly assigned to it, and the count has to be divided by the right denominator before it means anything.

RULE **An occupational count without an occupational denominator is a workload figure, not a risk figure.** It tells you how many deaths fell in this category. It cannot tell you whether people in that occupation are at higher risk than anyone else, because that comparison requires dividing by the number of people in the occupation, which is a different data source entirely.

14.2 The number, in full

The current published figure is this. **In 2024, 10,546 persons involved in the farming sector died by suicide in India, consisting of 4,633 farmers or cultivators and 5,913 agricultural labourers, accounting for 6.2 per cent of all suicide victims**. The denominator of that percentage is 1,70,746 total suicide victims recorded in India in 2024. The source is the National Crime Records Bureau's annual compilation of accidental deaths and suicides, the reference year is 2024, and the population is persons recorded by NCRB as involved in the farming sector, comprising the farmer or cultivator and agricultural labourer categories.

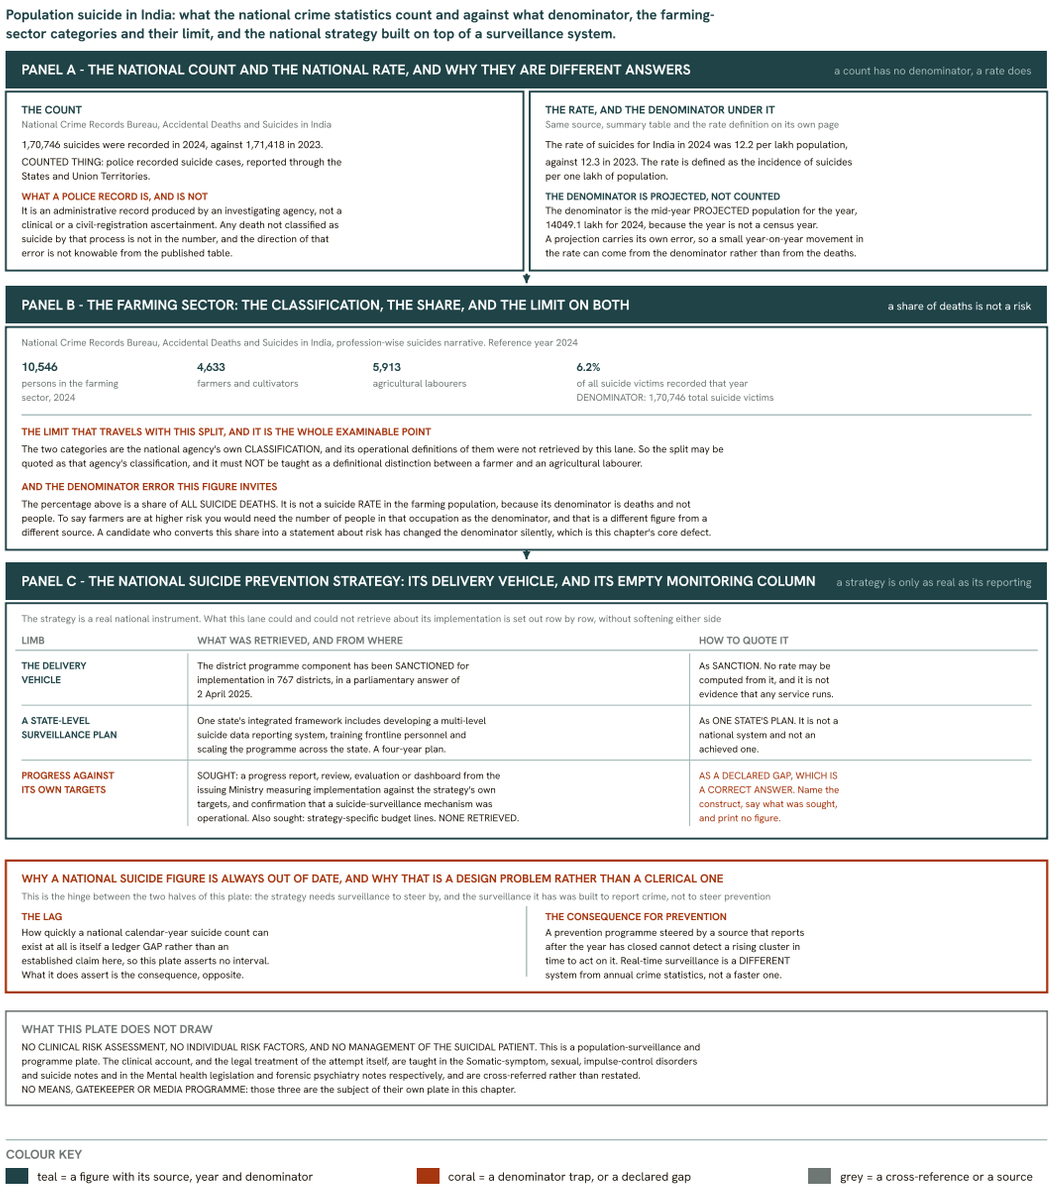


Fig 39.5 Population suicide in India: the count, the rate, the farming-sector categories and the strategy above them. The national agency recorded 1,70,746 suicides in 2024, against 1,71,418 in 2023, at a rate of 12.2 per lakh against 12.3, computed on a projected mid-year population because the year is not a census year. In the farming sector 10,546 deaths were recorded, comprising 4,633 farmers and cultivators and 5,913 agricultural labourers, 6.2 per cent of all suicide victims that year; because the agency's own operational definitions of those two categories were not retrieved, the split is quotable as its classification and not as a definitional distinction, and the percentage is a share of deaths rather than a rate in the farming population. The strategy's delivery vehicle is sanctioned in 767 districts, while progress against its own targets, confirmation that a surveillance mechanism was operational and its specific budget lines were all sought and none retrieved.

The accompanying figure sets out the surveillance architecture this figure comes out of, the categories the bureau publishes and the denominators available for each, alongside the governance of the national prevention strategy. Read it with the two limitations of the next two subtopics in mind, because the figure shows what the system counts and this text explains what the counting cannot support.

10,546	4,633	5,913	6.2%
FARMING-SECTOR SUICIDES RECORDED BY NCRB IN 2024	CLASSIFIED AS FARMERS OR CULTIVATORS	CLASSIFIED AS AGRICULTURAL LABOURERS	OF ALL RECORDED SUICIDE VICTIMS, NOT OF PEOPLE IN FARMING

Two features of that set deserve immediate attention, because both are routinely lost. The larger of the two component groups is agricultural labourers rather than landowning cultivators, which cuts against the way the issue is usually framed in public discussion. And the percentage is a share of all suicides in the country, which is a statement about the composition of the national total and not about risk in the farming population.

14.3 The classification, and the definition that could not be obtained

The split between the two categories is the most quoted feature of this data and the one this chapter is least able to explain. The bureau publishes the farming sector total divided into farmers or cultivators and agricultural labourers, and this chapter prints that split because it is what the source prints. What could not be obtained is the bureau's own operational definition of either category: the rule a recording officer applies to decide which of the two a particular deceased person belongs in.

That absence sets a hard limit on what may be said, and the limit is worth stating as a rule of method. The split may be reported as the bureau's classification. It may not be taught as a definitional distinction between two groups, because exactly what separates them in the bureau's coding is the thing that was not obtained. Anyone who tells you that the categories divide people by landholding, by tenancy status or by whether the work is paid is supplying an operational definition from inference, and the inference may be right, but it is not sourced.

This matters practically and not just pedantically. If a service or a state government wants to target the larger group, it needs to know who is in it. A category whose boundary rule is unpublished cannot be used to design an intervention with any confidence about whom it will reach.

14.4 The denominator problem, which is the heart of this topic

Here is the central analytical point of the section, and the one most worth carrying into an examination. To say that people in farming are at elevated risk of suicide, you need a rate: deaths in the farming population divided by the size of the farming population, compared with the corresponding rate elsewhere. The numerator exists and is published. The **denominator** does not exist in any form that can be validly matched to it.

The reason is that the two quantities are produced by different systems that do not share a definition of who is a farmer.

- **The numerator is an administrative classification made at the time of death,** by a recording officer, using a coding rule that this chapter could not obtain, applied to whatever information about the deceased was available at that moment.
- **Any candidate denominator is a census or survey classification,** made under its own definition of agricultural work, in a different reference period, and with its own treatment of people who do several kinds of work.
- **Neither system was designed to be divided by the other.** Combining them produces a rate whose numerator and denominator describe overlapping but non-identical populations, which is the classic route to a fabricated statistic that looks entirely respectable.

So the published share, being of all suicide victims rather than of people in farming, is doing something quite different from what most readers assume. It tells you what proportion of India's recorded suicides occurred among people classified as being in farming. It cannot tell you whether being in farming raises the risk, and it will move if the total number of suicides in the country changes even when nothing whatever happens in agriculture.

TRAP **Converting the share of all suicides into a statement about risk in farmers.** A share of the national total has the national total as its denominator, so it responds to changes anywhere in the country. Saying that farming accounts for a given percentage of suicides is correct. Saying that farmers are therefore at that level of risk, or comparing the share with the proportion of the workforce in agriculture as though the comparison were a risk ratio, is an error of denominators dressed up as an analysis.

IN INDIA

If you are asked to design a district response, do not begin from the national share. 6.2 per cent of all recorded suicide victims is a composition figure whose denominator is 1,70,746 total suicide victims recorded in India in 2024, and it tells you nothing about your district's farming population. Begin instead by establishing what your local denominator would be and who defines it, because until that is settled you cannot tell whether an intervention has worked, only whether the count went up or down for reasons you cannot attribute.

14.5 Who gets counted as a farmer, and who does not

Even before the denominator problem, the numerator has an assignment problem. Occupational classification at the time of death rests on what somebody records about a person's work, and agricultural work in India is unusually hard to record cleanly.

PERSON	HOW THE DEATH MAY BE CLASSIFIED	WHY THE ASSIGNMENT IS UNCERTAIN
An owner-cultivator working land held in their own name	Most likely within the farming sector	The clearest case, and the one public discussion imagines as typical
A tenant or sharecropper working land held by somebody else	Could fall in either component, or outside the sector	Whether cultivation without ownership counts as cultivating depends on a coding rule that is not published
A woman working on family land without holding it	Frequently recorded by household role rather than by occupation	Unpaid family labour is systematically under-recorded as work in many administrative systems
A person who farms in one season and does other paid work in another	Depends on which occupation the record captures	Multiple occupation is common and the record usually accommodates one
A landless labourer working across several holdings	Most likely the agricultural labourer component	Casual and migratory work is the hardest to attribute to any single occupation

Read that table as a description of measurement error with a direction rather than as a list of exceptions. The cases that are hardest to classify are concentrated among the poorest, the landless, the migratory and among women. Measurement error that is heavier at one end of a social gradient does not simply add noise; it biases the picture, and it biases it towards the group that is easiest to identify, which is the landholding cultivator.

14.6 What the data can and cannot say about causes

Determinant claims about this population are made constantly and are rarely qualified. Two limits govern how far a police-recorded surveillance system can support them.

The first is that a cause field in an administrative record captures the reason recorded by an investigator, from information available shortly after a death, in a context where the family may have strong reasons to present some explanations and not others. It is a recorded attribution, not an aetiological finding, and treating it as one imports the recorder's framework into what looks like epidemiology.

The second is that the same denominator problem prevents any comparison of determinant frequency between the farming population and everybody else. Without a rate, a determinant that appears often in this group cannot be shown to appear more often than in another group, which is what a risk factor claim requires.

The determinants of agrarian distress themselves, and the general aetiological models of suicide within which they sit, are taught in the Somatic-symptom, sexual, impulse-control disorders and suicide notes, and are not re-established here. What this section adds is the measurement caution that has to travel with them.

14.7 Prevention, and what would count as evidence

Targeted prevention for this population has been attempted in several forms, and this chapter holds no evaluation of any of them: no controlled comparison, no before-and-after estimate with a stable denominator, and no published attribution of change in the count to any programme. That absence is worth teaching directly, because it lets a candidate say something more useful than a list of interventions.

Set out what an evaluation would need, and the difficulty becomes obvious.

- **A stable denominator** over the evaluation period, so that a falling count can be distinguished from a shrinking population working in agriculture.
- **A comparison group** not exposed to the intervention, since counts in this area move with agricultural, economic and reporting conditions that no programme controls.
- **A consistent classification rule** across the period, because a change in coding practice would move the count on its own.
- **A pre-specified outcome and period,** since with year-to-year variation a favourable window can be found after the fact in almost any series.

Population-level prevention levers that apply here, including restriction of access to means and the national strategy that organises them, are taught in their own sections of this chapter, and this section does not restate them.

MNEMONIC

Who counted, who was counted, out of what

Three questions for any occupational suicide statistic. Which system produced the count, which coding rule assigned people into the category, and which denominator the figure is a proportion of. For this statistic the first is answerable, the second was not obtained, and the third does not exist in matchable form.

PITFALL

Comparing this year's count with last year's and calling the difference a trend. Year-on-year movement in an administratively recorded count reflects changes in reporting and classification practice as well as changes in deaths, and neither is separable without a stable denominator and a stable coding rule. Report the count with its year and its source; say that a trend claim requires both of those to have held constant, and that this chapter cannot establish that they did.

The farming-sector figure is a count with a national denominator, not a rate: 10,546 recorded farming-sector suicides in 2024, being 6.2 per cent of 1,70,746 total recorded suicide victims, and no valid occupational denominator exists to turn it into a risk.

15 · National suicide prevention strategy and population-level suicide prevention approaches

A national strategy is a governance document, not a treatment. It states what a country intends to do about a cause of death, names the sectors that will do it, and sets targets against which somebody is later supposed to measure it. That last clause is where most of the examinable content of this topic lives, because it is the clause that most often goes unhonoured. Candidates are rarely asked to recite a strategy's contents; they are asked what a strategy is for, what has to exist underneath it before it can work, and how one would know whether it had. Those are service questions, and this section answers them.

The prevention framework itself, the tiering of interventions into universal, selective and indicated, the four evidence-based population interventions in the World Health Organization's implementation guidance, the strategy's own headline goal and its staged action horizons, is taught in full in the Somatic-symptom, sexual, impulse-control disorders and suicide notes and is not restated here. What this section adds is the layer above and below that framework: how a strategy is delivered through an existing programme, what a surveillance system has to be able to do before a target can be checked against it, what the financing question is, and, unusually for a textbook section, a candid account of what is not published at all.

15.1 What a population-level strategy actually is, and what it is not

India has a **National Suicide Prevention Strategy**, which is the national governance instrument for suicide prevention and the reference point for every state and district plan written after it. A resident should be able to say three things about any such document. First, it is **multisectoral** by necessity: the levers that move population suicide mortality sit largely outside health, in agriculture, in commerce, in education and in the media, and a strategy that binds only the health ministry has already conceded most of its own case. Second, it is *staged*, because the levers act on different timescales; a regulatory restriction can act within a year, while a schools programme acts over a generation. Third, and this is the part candidates skip, it is *accountable*, meaning it names targets and therefore commits somebody to measuring them.

What a strategy is not is a service. It does not by itself open a clinic, train a worker or restrict access to anything. Everything it promises must be delivered through machinery that already exists or must be built, and in India that machinery is overwhelmingly the district-level mental health programme sitting under the national programme. This is why a question about the national strategy is very often, in substance, a question about district implementation.

RULE **A strategy is judged by its delivery vehicle and its measurement system, never by its text.** Two countries with identical strategy documents will produce different mortality curves if one has a functioning district service and a timely surveillance system and the other has neither. When you are asked to appraise a national strategy, appraise those two things first and the document second.

15.2 The delivery vehicle, and the single most misread figure in Indian community psychiatry

The vehicle for population suicide prevention in India is the **District Mental Health Programme**, the district-level component of the national programme. Its administrative scale is published: the district programme has been sanctioned for implementation in **767** districts. That is a governmental statement of approval, made in reply to a question in Parliament, and it is the number a candidate will meet on a slide.

Read what it counts. It counts districts for which implementation has been *sanctioned*. It does not count districts in which a suicide-prevention service exists, in which a gatekeeper has been trained, or in which a person in crisis could be seen. Sanction is an act of the centre; function is a property of the district. The two are separated by staffing, by the release of funds, by the state's own absorption capacity and by time, and nothing in a sanction figure carries information about any of them. Nor does the figure support a rate: there is no valid denominator against which a count of approved districts becomes a proportion of anything meaningful, because approval is cumulative and administrative rather than a measurement of a population at a point in time.

TRAP **Reading a sanctioned-district total as coverage.** Candidates write that the programme "covers" the sanctioned number of districts and then reason from that to national access. The word in the source is approval, and the correct sentence in an answer is that the programme is sanctioned in that many districts while the number in which suicide-prevention services actually function is not established. That second half of the sentence is worth more marks than the number is, because it shows you know what the number measures.

IN INDIA

When you are asked what a district psychiatrist can actually do under the national strategy, answer through the district programme, because that is the only route the strategy has. The practical consequence is that a district's suicide-prevention capability is decided by whether its programme post is filled and its funds released, not by whether the district appears on the sanctioned list. A resident asked to plan district suicide prevention should begin by establishing which of the sanctioned components are staffed and running locally, and treat the national total as a statement about paperwork rather than about services.

15.3 Surveillance read as an instrument, not as a table of numbers

A strategy that sets a mortality target needs a measurement system capable of telling it whether the target moved. In India **suicide surveillance**, meaning the routine national counting of deaths by suicide for programme purposes, is produced by the police-reporting system and published annually as a compendium of accidental deaths and suicides. Treated as a data instrument rather than as a source of quotable totals, it has three properties that decide what it can and cannot be used for.

- **Case definition.** The unit is a case recorded by police, which is an administrative event with legal consequences and not a clinical determination. Anything that changes the willingness of families or of the police to record a death as suicide changes the series without anything changing in the population.

- **Timing.** The annual publication collates a full calendar year, so collection can only begin once that year has ended. The interval between an event and its appearance in the national series is therefore at least a year and in practice longer.
- **Granularity.** The series is compiled upward from states, so its usefulness for targeting depends on the quality of state-level reporting, which is not uniform.

The timing property is the one with programme consequences and it is worth stating plainly as a declared limitation rather than glossing it. A system whose national picture arrives a year or more after the events it describes cannot function as an early-warning instrument, cannot support the rapid evaluation of a newly introduced intervention, and cannot tell a state within a planning cycle whether what it did last year worked. Whether India now has a suicide-surveillance mechanism that is operational, nationally complete and timely enough to monitor strategy outcomes is not established: this chapter sought an official assessment of exactly that question through the health ministry, the bureau that publishes the national series and the national institute that advises on the programme, and none was retrieved. The honest examination answer is that the strategy names a surveillance objective and that no published assessment confirms the objective has been met.

Sanction	None	Over a year	Not established
WHAT THE DISTRICT TOTAL COUNTS, NOT SERVICE DELIVERY	VALID DENOMINATOR FOR THE DISTRICT TOTAL, SO NO RATE	MINIMUM LAG BEFORE A CALENDAR YEAR ENTERS THE NATIONAL SERIES	PROGRESS MEASURED AGAINST THE STRATEGY'S OWN TARGETS

15.4 The state as the real implementing unit

Health is administered by the states, so a national strategy is realised as a set of state strategies and state programmes, and the interesting variation is between states rather than around the national mean. At least one state has taken this seriously enough to build an integrated framework of its own for suicide prevention and response, spanning health workers, communities and law-enforcement agencies. Its programme is designed around three components that are worth learning as a template: **a multi-level suicide data reporting system, training for frontline personnel, and scale-up of the programme across the whole state**.

Notice the order and the logic. The reporting system comes first because without it the state cannot know its own baseline or detect change; training comes second because it converts the strategy into an act somebody performs; scale-up comes third because a pilot in two districts answers no population question. That sequence is the generalisable teaching point, and it is a better answer to "how would you implement the national strategy in your state" than any recitation of the national document.

The specifics of any individual state's instrument, including the date on which such a strategy was released and the contents of its published text, are not established in this chapter's sources and should not be reconstructed from a seminar handout. Name the state layer, describe the template, and say that the state documents themselves are the primary sources a candidate would need to read.

15.5 Financing, which is where most strategies are actually decided

A strategy without money attached is a statement of intent. This is not a cynical observation; it is the standard first question in any appraisal of a public-health programme, and it has been raised officially in India in relation to both the mental health legislation and the suicide prevention strategy, where the absence of dedicated budgetary allocation has been identified as a concern about implementation. An early official parliamentary assessment recorded that whether the strategy was being implemented effectively remained undetermined.

What this chapter cannot tell you, and says so rather than inventing it, is how much money the strategy has. Allocations, releases and utilisation specific to the suicide prevention strategy, separated from the broader mental health programme and from the national institutions, were sought from the health ministry, from budget documentation and from parliamentary answers, and were not retrieved. Two things follow for an examination answer. First, do not quote a rupee figure for the strategy; there is none you can defend. Second, the fact that a national strategy's own financing cannot be isolated from the wider programme is itself a finding about the strategy, and saying so demonstrates more understanding than a number would.

PITFALL

Treating an allocation as an outcome. Where a programme's spending is published, clinicians frequently quote the allocation as evidence of activity. An **allocation** is a budget line, a **release** is money that has left the treasury, and **utilisation** is money that has been spent; the three routinely differ, and none of them is a service. If you use a financial figure at all, say which of the three it is, and for which year.

15.6 Evaluating a strategy against its own targets

The cleanest way to appraise any national strategy is to hold it to the standard it set for itself, because that standard was chosen by the authority that wrote it and cannot be dismissed as an outsider's benchmark. That approach requires a progress report, a review, an evaluation or a live dashboard that measures implementation against the strategy's stated targets, published by the issuing body or by an independent auditor.

No such assessment was retrieved for the Indian strategy. The search covered the health ministry's own programme pages and its published strategy location, parliamentary answers in both houses, a parliamentary standing committee report, government press releases and the national institute's own domain. The absence is the finding, and it should be taught as one: India has a national strategy, a delivery vehicle with a published sanction footprint, a surveillance system with a structural reporting lag, and no published measurement of the strategy against its own targets. A candidate who can lay out those four statements in order has given a better account of Indian suicide prevention than one who recites the strategy's contents.

It is worth setting the four appraisal questions out against each other, because the pattern that emerges is the section's real teaching point. In every case the design statement is published and traceable to an authority, and in every case the delivery statement is either missing or is a different quantity wearing the design statement's clothes.

APPRAISAL QUESTION	WHAT IS PUBLISHED AS DESIGN	WHAT DELIVERY WOULD REQUIRE	STATE OF THE DELIVERY EVIDENCE
Reach	Districts for which the district programme is sanctioned, stated by the health administration in reply to Parliament	Districts in which a suicide-prevention service is staffed and seeing people	Not established; sanction and function are different quantities and only the first is published
Measurement	A surveillance objective inside the strategy, and an annual national compendium of recorded suicides	A system that is nationally complete and timely enough to detect change within a planning cycle	Not established; the annual series can only begin collecting a calendar year once it has closed
Money	The mental health programme's own budget lines	Allocation, release and utilisation attributable to the strategy itself	Not retrieved; strategy-specific financing cannot be separated from the wider programme
Effect	Targets stated by the strategy for itself	A progress report or evaluation from the issuing body or an auditor, measured against those targets	Not retrieved from the ministry, from Parliament, from a standing committee or from the national institute

Two features of that table are examinable in their own right. The first is that the left-hand column is answerable for almost any country and the right-hand column is answerable for very few, which is why international comparison of suicide-prevention strategies is so much weaker than it looks. The second is that a candidate who has learned only the left-hand column will answer a delivery question with a design statement and will not notice the substitution, because the design statement is a real, sourced, governmental figure. The error is not one of accuracy but of scope, and scope errors are invisible from the inside.

MNEMONIC

Vehicle, Surveillance, Sectors, Money, Measurement

The five questions to ask of any national suicide prevention strategy, in the order in which they decide whether it can work. What delivers it, what measures it, who outside health is bound by it, what funds it, and who checks it. Every one of the five is answerable for India, and two of the answers are that nothing is published.

15.7 Writing about this in an examination

Three habits separate a strong answer from an average one here. The first is scope discipline: every national figure you print must travel with the body that issued it, the period it refers to, the population it describes and what it is a proportion of, and if you cannot supply all four you should describe the thing without its number. The second is the distinction between design and delivery, which recurs in every subtopic above: a sanctioned district, a stated objective and an allocated budget are all statements of design, while a functioning service, a measured outcome

and spent money are delivery. The third is the willingness to declare an absence. In a viva, "that figure is not reliably published, and here is what was looked for" is a stronger answer than a confidently quoted number of unknown provenance, and it is the answer this chapter is built to let you give.

The other suicide-prevention material in this chapter divides cleanly. The specific population interventions, means policy, gatekeeper programmes and media guidance, are configured and evaluated in their own section; the agrarian category argument and its denominator problem belong with farmer suicide; the national data series read as a distribution belongs with suicide epidemiology. The clinical prevention framework, the intervention tiers and the international implementation guidance behind all of them belong to the Somatic-symptom, sexual, impulse-control disorders and suicide notes.

> A national suicide prevention strategy is only as real as its delivery vehicle, its surveillance timeliness and its financing, and in India the first is sanctioned in a published number of districts, the second reports a calendar year only after that year has closed, and the third cannot be separated out at all.

16 · Means restriction, gatekeeper training and media guidelines in suicide prevention

Three interventions, and each of them is two different things. Means restriction, gatekeeper training and media guidance can each be described as something a clinician does with a person in front of them, and each can be described as something a state designs, funds, delivers and evaluates across a population. The clinical readings are taught elsewhere and are cross-referred below. This section takes the second reading throughout, because that is where the marks in a public-health paper sit and because the two are genuinely different objects: a gatekeeper programme is commissioned and evaluated, a gatekeeper referral is made.

The mechanisms behind the three, why restricting access to a method works when crises are short and ambivalent, what the evidence says about gatekeeper training's effect on recognition and on mortality, and the paired imitation and protective effects of media coverage, are taught in full in the Somatic-symptom, sexual, impulse-control disorders and suicide notes. Means restriction as a task carried out with a young person and their family, including the household conversation and the safety plan it belongs to, is taught in the Child behavioural, emotional and other disorders notes. Neither is restated here. What follows is the Indian implementation layer: which instruments exist, who issues them, what has actually been measured, and what has not.

16.1 The three levers, read as programmes

A population lever has four parts that a clinical act does not: an issuing authority, a legal or administrative instrument, a delivery route, and an evaluation endpoint. Run the three interventions through those four parts and the differences between them become obvious.

- **Means restriction** is issued by a regulator, not by a health department, and its instrument is usually a statute or a subordinate order. Its delivery route is the market and the supply chain

rather than the clinic, which is why it reaches people who never present to a service. Its endpoint is method-specific mortality.

- **Gatekeeper training** is issued by a health or education authority, its instrument is a curriculum, its delivery route is a workforce, and its endpoint is contested: recognition and referral are measurable, whereas an effect on deaths is not demonstrable from a training programme alone.
- **Media guidance** is issued by a press or broadcast regulator, its instrument is a code of conduct, its delivery route is editorial practice, and its endpoint is the content of reporting itself rather than any clinical event.

That table of four parts is the frame for everything below, and it is also the answer to the commonest structural mistake in this topic, which is to compare the three on effectiveness as though they were competing treatments. They are not substitutes. They act on different populations through different institutions, and a national programme runs all three because each reaches people the other two miss.

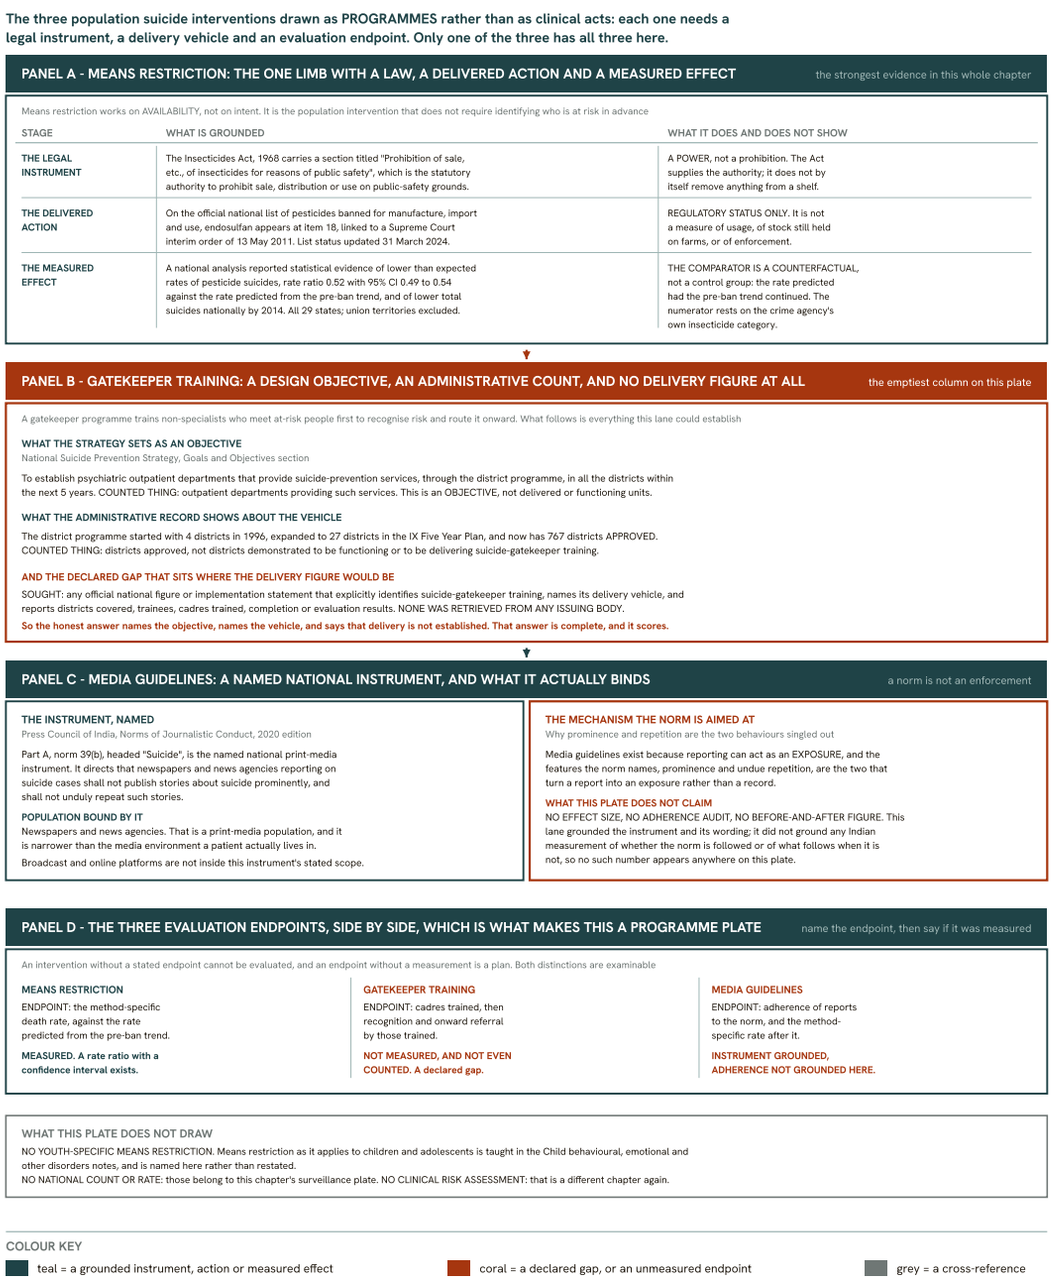


Fig 39.6 Means restriction, gatekeeper training and media guidelines, each drawn as a programme with its own evaluation endpoint. Means restriction is the one limb with all three stages grounded: a statutory power to prohibit an insecticide for reasons of public safety, a delivered ban listing endosulfan at item 18, linked to a Supreme Court interim order of 13 May 2011, and a measured effect, a pesticide-suicide rate ratio of 0.52 with 95% CI 0.49 to 0.54 against the rate predicted from the pre-ban trend. Gatekeeper training has an objective set by the national strategy and a vehicle approved in 767 districts, but no official figure naming trainees, cadres, completion or evaluation was retrieved from any issuing body. The media limb has a named national instrument directing that newspapers shall not publish suicide stories prominently or unduly repeat them, bounded to print media, with no Indian adherence measurement grounded here.

RULE **Name the issuing authority before you discuss the intervention.** A population lever that no institution is empowered to pull is a wish. In India the three levers sit with three different authorities, only one of which is the health ministry, and half of the implementation difficulty in this topic follows from that single fact.

16.2 Means restriction in India: which instruments actually exist

India regulates agricultural chemicals through a dedicated statute. The **Insecticides Act, 1968** carries a provision, at Section 27, headed as a prohibition on the sale and related dealings in an insecticide for reasons of public safety. Read as a public-health instrument rather than as an agricultural one, that is the legal hook on which access restriction hangs: it is the power to take a compound off the market, and it exists independently of any mental health law.

The operational expression of that power is a published national list of pesticides that are banned, refused registration or restricted in use, maintained by the plant-protection directorate. Compounds move onto it by different routes, including judicial ones: **endosulfan appears in the banned-for-manufacture, import and use section of that list at item 18, linked to an interim order of the Supreme Court made in 2011**. That detail is worth carrying because it shows the mechanism by which access restriction is achieved in practice, which is administrative and judicial rather than clinical, and because it locates the document a resident should cite.

It is worth pausing on why the instrument is agricultural rather than medical, because the point generalises. Population access restriction almost always runs through the law that governs the commodity, not through the law that governs health: a pesticide through an insecticides statute, a medicine through drug scheduling and pack regulation, a structure through building and transport rules. The health system's role in each case is to supply the mortality evidence that justifies the restriction and then to measure what follows, not to issue the restriction itself. A resident who understands that will not waste an answer proposing that a mental health authority ban anything, and will instead describe the correct route, which is evidence to the sectoral regulator and, where that fails, to the courts.

Two cautions travel with any statement about that list. It records *regulatory status*, not use and not enforcement; a compound can be banned on paper and remain in circulation in a district. And a list is a statement about compounds, not about deaths, so nothing about mortality follows from the list itself. The mortality question is separate and is answered below.

16.3 What the regulation bought, and how anyone knows

Regulation of a method is one of the few population interventions in psychiatry that can be evaluated at national scale, because the exposure changes on a known date for the whole country and the outcome is recorded routinely. The Indian evaluation used exactly that design. Comparing the observed rate of pesticide suicides against the rate predicted by extrapolating the stable pre-ban trend, a national analysis reported **a pesticide-suicide rate ratio of 0.52, with a 95 percent confidence interval of 0.49 to 0.54**, together with a finding that total suicides nationally were lower than expected by 2014.

Learn the structure of that claim, not just its value. The comparator is a counterfactual, a modelled rate rather than an observed one, which means the estimate inherits every assumption in the trend model. The outcome is drawn from routine national records of deaths classified as pesticide poisoning, and the way those record categories are defined, and what they can and cannot bear, is the argument taken up where farmer suicide is taught in this chapter; it is not repeated here. The second finding, that total suicides also fell below the predicted line, is the analytically interesting one, because it is the finding that speaks to whether restriction merely displaces the method.

PITFALL

Quoting a rate ratio as though it were a proportion. A ratio below unity from a counterfactual analysis says the observed rate was about half of what the pre-ban trend predicted for that year. It does not say half of suicides were prevented, it does not say half of pesticide deaths stopped, and it cannot be converted into a number of lives without the predicted rate and the population it applies to. State it as a ratio against a predicted rate, or do not state it at all.

16.4 Gatekeeper programmes: what has to be specified before one exists

A **gatekeeper** programme is a training intervention aimed at people who meet those at risk before any clinician does. As a programme rather than as a course, it is defined by six specifications, and a plan that omits any of them is not yet a programme: which cadre is trained, what competency the training is supposed to produce, who delivers and refreshes the training, what the trainee is expected to do when they identify somebody, what receives the referral, and what will be counted afterwards. The sixth is the one most often left blank, and it is the one that decides whether the programme can ever be shown to have worked.

India's obvious delivery routes for such a programme are the district mental health programme and the national tele-counselling service, both of which already reach a workforce and a public. What is *not* available is any official national statement that identifies suicide-gatekeeper training as a named component, names its delivery vehicle, and reports districts covered, personnel trained, cadres included, completion rates or evaluation results. That was sought from the health ministry, the health services directorate, the national health mission and the bodies responsible for programme implementation, and it was not retrieved. Say so in an answer, because it is the truthful position and because it explains why the Indian literature on gatekeeper training consists of individual studies rather than programme reports.

Choosing that endpoint honestly is harder than it looks. Recognition and referral behaviour are the endpoints a training programme can plausibly move and can measure within a reasonable period, and they are legitimate. Deaths are the endpoint everybody wants, and a single training programme is almost never powered to detect a change in them, because the event is rare in any one district and because a trained workforce is only one link in a chain that also requires a service to refer into. Writing mortality into the evaluation plan of a district training programme is therefore not ambition; it is a design error that guarantees a null result and discredits a

reasonable intervention. State the intermediate endpoint, state that it is intermediate, and say what would have to be true downstream for it to matter.

TRAP **Treating an approved district as a trained district.** The national administrative scale of the district programme is published, and it is regularly imported into answers about gatekeeper coverage. Approval of a district under the programme carries no information at all about whether anybody in that district has been trained to recognise suicide risk. Approval is an accounting fact about the centre; training is a fact about a workforce, and no published national count of the second exists.

16.5 Media guidance: the instrument, and the reach it does not have

A **media guideline** is a code binding on publishers, and in India the named national instrument for print is the Press Council of India's **Norms of Journalistic Conduct, 2020**, which contains a suicide-specific norm directed against reporting suicide prominently and repeating such stories unduly. Two features of that instrument matter more than its contents.

The first is its scope. It is a code for newspapers and news agencies, which is to say for print journalism and the agencies that feed it. Broadcast media and the online and social platforms on which most people now encounter a suicide story sit under different regulators or under none, so compliance with a print code is not compliance across the media environment. The second is its force. A council code operates by adjudication and censure rather than by penalty, so its effect depends on editorial willingness in a way that a pesticide ban does not.

The practical consequence for a resident is that the evaluation endpoint for media guidance is the content of coverage, measured against the code's own requirements, rather than any mortality outcome. If you are asked how you would evaluate an Indian media-guidance intervention, the answer is a content audit of reporting before and after, sampled across the outlets the code actually binds, and an honest statement that the outlets it does not bind are the ones carrying most of the audience.

16.6 Delivering all three through one district system

Whatever the national strategy says, everything it promises arrives in a district through a single service. The strategy addresses this directly: **the 2022 strategy sets, as a design objective, the establishment of psychiatric outpatient departments providing suicide-prevention services through the district mental health programme in all districts within the next five years**. That is a target, and reading it as an achievement is the single commonest error in this topic.

Against that target sits the administrative position: the health services directorate reports the district programme's national scale as **767** approved districts. Approval is what that number counts. It is not a count of districts with a functioning suicide-prevention outpatient department, and it must not be rewritten as gatekeeper delivery or as any other service output. The gap between a design objective expressed in all districts and an approval figure expressed in a count of districts is not a contradiction; the two are simply different quantities, and an answer that treats them as one has made a scope error rather than an arithmetic one.

Regulator	Press council	District programme	Not established
ISSUING AUTHORITY FOR MEANS RESTRICTION IN INDIA	ISSUING AUTHORITY FOR THE PRINT MEDIA CODE	DELIVERY ROUTE THE STRATEGY NAMES FOR ALL DISTRICTS	NATIONAL GATEKEEPER TRAINING COVERAGE, ANY CADRE

16.7 The evidence state of each lever, side by side

The most useful thing a candidate can carry out of this section is not a set of figures but a map of what is and is not established for each lever in India. The pattern is asymmetric in an instructive way: the lever with the strongest legal instrument is also the only one with a national effect estimate, and the lever that is easiest to run is the one with no national data at all.

LEVER	INDIAN INSTRUMENT, AND WHO ISSUES IT	WHAT IS ESTABLISHED ABOUT EFFECT	WHAT IS DECLARED NOT ESTABLISHED
Access restriction	A dedicated insecticides statute with a public-safety prohibition power, plus a national banned and restricted list maintained by the plant-protection directorate	A national counterfactual analysis reporting a pesticide-suicide rate ratio below unity against the predicted pre-ban trend, with a total-suicide finding in the same direction	Enforcement and actual availability at district level, which regulatory status does not measure
Gatekeeper programmes	No named national instrument retrieved; the district programme and the national tele-counselling service are the available delivery routes	Nothing at national level	Districts covered, personnel trained, cadres included, completion and any evaluation result
Media guidance	A press council code of journalistic conduct with a suicide-specific norm, binding on newspapers and news agencies	Nothing at national level as an outcome; the endpoint is compliance in content	Compliance auditing, and coverage of broadcast and online outlets the code does not bind

IN INDIA

The operational decision this section forces is which lever a district team can actually pull. Access restriction is not theirs: it sits with an agricultural regulator and with the courts, so the district contribution is surveillance and advocacy rather than action. Media guidance is not theirs either, beyond a local relationship with editors. Gatekeeper training is the one lever a district service can commission itself, which is precisely why it is the lever most often run without a specified endpoint. If you are asked to plan district suicide prevention, commission the training you control, and write the count you will report before the first session is delivered.

MNEMONIC

Authority, Instrument, Route, Endpoint

The four parts of any population lever, and the order in which to write them. Missing authority makes the intervention unenforceable; missing instrument makes it voluntary; missing route makes it undeliverable; missing endpoint makes it unevaluable, which is the failure mode of gatekeeper programmes worldwide.

The remaining suicide material in this chapter divides as follows. The national strategy as a governance instrument, its delivery vehicle and its financing are taken up in their own section; the national data series read as a distribution belongs with suicide epidemiology; the occupational category argument belongs with farmer suicide. The mechanisms and the international evidence base for all three levers remain with the Somatic-symptom, sexual, impulse-control disorders and suicide notes, and the family-level task with the Child behavioural, emotional and other disorders notes.

Means restriction is the only one of the three levers with an Indian legal instrument and an Indian national effect estimate; gatekeeper training is the only one a district can commission for itself and has no national coverage figure at all; media guidance binds print alone and is evaluated on the content of reporting, never on deaths.

17 · Epidemiology of suicide in India and high-risk groups

Every national suicide figure you will ever quote is the output of a data system, not a measurement of the world. That is not scepticism, it is the first thing a public-health reader has to understand about this topic, and it is what separates an answer that recites a number from an answer that can defend one. India's national suicide statistics come from a police-reporting chain, are published annually, and describe a case category with a legal existence. Everything that follows about counts, rates, trends and high-risk groups is conditioned by that fact.

The clinical material adjacent to this section is owned elsewhere and is not restated. Individual risk assessment, the weighing of risk against protective factors, structured evaluation and the limits of prediction belong to the Somatic-symptom, sexual, impulse-control disorders and suicide notes, as do the aetiological frameworks that explain why populations differ. Custodial suicide prevention, including screening on reception into prison and the advisory machinery around it, belongs to the Forensic ethics, consent and legal procedure notes. What this section owns is the surveillance instrument itself, the distribution it produces, and what that distribution licenses a programme to do.

17.1 What the national statistic is a statistic of

The published national series is compiled by the **National Crime Records Bureau, the NCRB**, from **police-recorded suicide cases** reported upward by states and union territories and collated for a calendar year. Three consequences follow immediately and each is examinable.

- The unit of count is a *recorded case*, so anything that changes recording changes the series. A family's reluctance to have a death recorded as suicide, a local practice of recording an ambiguous death as accidental, or a change in police procedure will all move the number without anything moving in the population.
- The determination is administrative rather than clinical. No psychological autopsy stands behind a national count, and the intent judgment is made in a legal rather than a diagnostic frame.
- The compilation is annual and closed, which means the whole apparatus is retrospective by at least a full year. A surveillance system with that latency describes the past accurately and cannot warn about the present.

RULE **Name the recording system before you quote its number.** "The suicide rate in India is X" is an incomplete sentence. The complete sentence is "the police-recorded suicide rate in India for the year Y, per lakh of the projected mid-year population, is X". Everything a candidate is likely to get wrong about this topic is contained in the words the short version leaves out.

17.2 The count and the rate, and how they differ

For the most recent published year the Bureau **recorded 1,70,746 suicides in India in 2024, against 1,71,418 in 2023**. The corresponding rate is defined as the incidence of suicides per one lakh of population, a lakh being 1,00,000 people, and the recorded rate is a series rather than a constant, so a figure quoted for an earlier reference year in the Somatic-symptom, sexual, impulse-control disorders and suicide notes belongs to an earlier point on the same series and does not compete with this one: for **2024 the rate is 12.2 per lakh, against 12.3 in 2023**.

A count and a rate answer different questions and are not interchangeable. The count answers a service-planning question, because a mortuary, a crisis service and a bereavement service all deal in people rather than in rates. The rate answers a comparison question, because only a rate can be set beside another year, another state or another country. A year in which the count falls while the population grows will show a falling rate for two different reasons at once, and separating them requires the denominator.

1,70,746	12.2	Recorded case	Over a year
POLICE-RECORDED SUICIDE CASES, INDIA, 2024	PER LAKH OF THE PROJECTED MID-YEAR POPULATION, 2024	THE UNIT BOTH FIGURES COUNT, NOT A CLINICAL DETERMINATION	LATENCY BEFORE A CLOSED CALENDAR YEAR IS PUBLISHED

17.3 The denominator, which is where rates are quietly built

India conducts a census at long intervals, so for every year that is not a census year the population **denominator** is not counted but projected. The published rate for the most recent year uses a projected mid-year population of 14049.1 lakh, and the rate for any non-census year is computed on a projection of that kind rather than on an enumeration.

That has two implications a resident should be able to state. The first is that a national rate carries the uncertainty of the projection as well as the uncertainty of the numerator, and neither is usually reported with an interval. The second is that projections drift: the further a year sits

from the last enumeration, the more the denominator depends on assumptions about fertility, mortality and migration, and a small systematic error in the denominator moves every rate in the series in the same direction. A trend computed across a long inter-census interval is therefore a weaker object than it looks, and this is worth saying explicitly in any answer that discusses whether the Indian rate is rising.

TRAP **Comparing a rate for one population against a count for another.** The classic version is to set a national rate beside a group's absolute number and conclude that the group is over-represented. A count of deaths in any group is a function of the size of the group, and without the group's own denominator it says nothing about risk. Where a valid denominator for a subgroup does not exist, the honest position is that the group's rate cannot be computed, not that its count is small or large.

17.4 What the system cannot see, and what else could see it

Undercount in a police-based series is not a defect to be apologised for; it is a structural property with a describable mechanism. Deaths that never reach police attention are absent; deaths that reach it but are recorded under another manner are misclassified; deaths recorded correctly but not transmitted upward are lost. Each of those acts in one direction, so the national series is best understood as a floor rather than an estimate with symmetric error.

The direction of that structural bias has a practical corollary that examiners like. If the series is a floor, then a district or a state reporting a low recorded rate has two possible explanations, a genuinely low incidence and a weak recording chain, and the second explanation is at least as common as the first. A programme that reads a low recorded rate as reassurance and withdraws attention from that district has therefore drawn exactly the wrong conclusion from a weak signal. The correct reading of an unexpectedly low figure is that the recording chain should be examined first and the population second.

India also maintains civil registration of deaths and a sample-based demographic survey, and the natural analytic move is to compare a police-derived suicide count against a registration-derived one for the same period. That comparison is the correct way to ask how large the undercount is. This chapter's sources do not establish the size or the direction of the difference for India, so no reconciliation figure is offered here, and a candidate should say that the comparison is the right method and that its published answer is what they would look for.

17.5 High-risk groups, read as population strata

A **high-risk group** in a public-health frame is not the same object as a high-risk patient. It is a stratum of the population whose rate exceeds the population rate, identified so that a finite intervention can be aimed somewhere rather than everywhere. Three things have to be true before a stratum is useful for targeting, and the national series satisfies only the first of them for most groups.

- **The stratum must be identifiable in the data.** The national record carries profile characteristics collected on the police form, so the series can be cut by sex, by age band, by state and by occupation and other social descriptors. The definitions of those recorded

categories, and what they will and will not bear, are the argument taken up where farmer suicide is taught in this chapter, and they are not repeated here.

- **The stratum must have its own denominator.** Without a population count for the group itself, a group rate cannot be computed and over-representation cannot be established. This is the point at which most confidently quoted group statistics fail.
- **The stratum must be reachable.** A group that can be described but not addressed through any existing service or regulatory route generates a finding, not a programme.

The third condition is worth dwelling on because it is where epidemiology becomes service planning. Strata differ enormously in how reachable they are. A group defined by an occupation can be reached through the sector that employs it and through the regulator that governs its inputs. A group defined by contact with a health service can be reached at that contact, which is why service-contact populations attract prevention effort out of proportion to their share of deaths. A group defined only by a demographic characteristic, such as an age band, has no institution attached to it at all, and reaching it requires a mass channel such as schooling, broadcast media or a helpline. When you are asked to design a targeted intervention, name the institution before you name the intervention, because the institution is what makes targeting possible.

Groups conventionally discussed in the Indian literature include young adults, students, people in agrarian occupations, people in custody and people with severe mental illness in contact with services. Naming a group is legitimate; attaching a proportion to it is not, unless the proportion carries its own denominator. This section prints no group-specific proportions, because none is established in its sources, and that absence should be stated in an answer rather than filled from a slide.

> **PITFALL**
>
> Confusing relative and absolute risk when planning. A stratum can carry a high relative risk and contribute few deaths because it is small, while a low-risk but very large stratum contributes most of the total. Targeting the first reduces risk for individuals; only action on the second moves the national count. A plan that names only high-relative-risk groups will be defensible clinically and useless at population scale, and this is the single most common weakness in resident-written prevention plans.

17.6 Inter-state variation, and why it is the useful unit

Health is administered by states, so the national rate is an average over administrative units that differ in almost everything that matters: recording practice, service availability, urbanisation, the dominant method available, and the age structure of the population. A national figure is therefore the least actionable number in the series. The state figure is the one a programme can be built on, and the district figure is the one a service can be built on.

Two cautions apply when reading variation between states. Differences in the recorded rate confound true differences in incidence with differences in recording, and there is no way to

separate them from the published series alone. And states differ in age structure, so a raw comparison between them is not a like-for-like comparison; the standard remedy is **age standardisation**, and any inter-state comparison that has not been standardised should be described as a comparison of **crude rates**. This chapter prints no state figures, because none is in its sources.

17.7 Reading the trend as a data series

The two most recent published years sit close together, with the count and the rate both marginally lower in the later year. A one-year movement of that size is not a trend and should never be described as one. A movement in a routinely collected series can arise from a real change in incidence, from a change in recording behaviour, from a change in the projected denominator, or from ordinary year-to-year variation, and a single comparison cannot distinguish them.

What an examiner is looking for is the discipline to say all four aloud, in that order, before offering an interpretation. If a candidate must commit, the defensible statement is that the national recorded rate has been broadly stable across the two most recent published years, that the level is what matters for planning rather than the direction of the last step, and that any claim about a sustained trend requires the series over many years together with a stated treatment of the denominator.

IN INDIA

The operational decision this section drives is what you plan against. Because the national series is retrospective by more than a year, is administratively defined, and averages over states that differ in recording as much as in incidence, a district service should not wait for the national publication to tell it what its problem is. It should treat the national figure as the background level and build its own local numerator from the sources it can see directly, meaning its own casualty and mortuary records and its own crisis presentations, while accepting that a locally built count is a count and never a rate unless the district's population denominator is stated with it.

17.8 What the distribution licenses, and what it does not

The final discipline is to be clear about what an epidemiological description entitles anyone to do. It licenses the aiming of finite resources at strata with a demonstrably higher rate, provided that the rate has a denominator. It licenses surveillance of the strata over time, which is the only way to know whether targeting worked. It does not license a prediction about any individual, and it does not license the inference that a correlate observed at population level operates as a cause in a person.

QUANTITY	WHAT IT ANSWERS	WHAT IT NEEDS TO BE VALID	WHAT IT CANNOT DO
National recorded count	How many recorded cases arose in a closed calendar year	A stated recording system and a stated year	Support comparison across years or places without a denominator
National recorded rate	How the level compares with another year, state or country	The same numerator plus a stated population base, projected for non-census years	Describe any subgroup, and cannot be applied to an individual
Group rate	Whether a stratum is over-represented	A numerator and a denominator for that same group	Be computed at all where the group's population base is unavailable
Trend	Whether the level is moving	Many years, a stable case definition and a stated denominator method	Be read from a single year-on-year step

India's national suicide figures are police-recorded counts and rates for a closed calendar year, computed for non-census years on a projected mid-year population, so they describe the past accurately, warn about nothing in the present, and support a statement about any subgroup only where that subgroup has a denominator of its own.

Stigma and anti-stigma interventions

18 · Stigma: concepts, types (public, self, structural) and measurement instruments

Stigma is not one thing, and a service that treats it as one thing will intervene in the wrong place. The word covers at least four distinct objects, each sitting in a different location, each moved by a different intervention and each measured by a different instrument. Typing stigma is therefore not an academic refinement; it is the step that decides whether an anti-stigma budget buys a media campaign, a workforce training programme, a change to an insurance rule or a support group for families. This section supplies the typology, its consequences for access to care, and the measurement layer that lets any of it be evaluated.

Two neighbouring accounts are owned elsewhere and are named rather than rebuilt. The social-psychological mechanism behind stigma, prejudice as a structure of belief, feeling and behaviour, the theoretical lineage that explains how in-groups and out-groups form, and the intergroup evidence base, belongs to the Social, Learning and Personality Psychology notes. The clinical reading, in which stigma is treated as a variable inside one consultation and divided into what a patient has met, what a patient anticipates and what a patient has internalised, with concealment as the behavioural pathway to a worse outcome, belongs to the Consultation-liaison, geriatric and transcultural psychiatry notes. What follows is the population and structural reading, which no other chapter carries.

18.1 Why a service types stigma rather than deploring it

Deploring stigma is free and changes nothing. A public-health reading asks four questions of any stigma phenomenon, and the answers differ by type.

- **Where does it live?** In a population, in a person, in a rule, or in a family.
- **Who holds it?** The answer identifies the respondent any later measurement will have to use.
- **What does it do?** Specifically, what does it do to whether a person seeks care, receives care and keeps a role, because a stigma that changes nothing is a sentiment rather than a public-health problem.
- **What would count as evidence that it had shifted?** This has to be answerable before the intervention is designed, not afterwards.

Running those four questions produces the typology below rather than presupposing it, which is the right order and the one an examiner rewards. Note that the fourth question is the one that most often has no good answer, and that a type of stigma for which no endpoint can be named is a type nobody will be able to show they reduced.

RULE **Locate the stigma before you choose the intervention.** An intervention aimed at the wrong location will produce a measurable change in the wrong outcome and no change in access. A campaign that improves public attitudes will not amend a discriminatory insurance rule, and a support group that reduces a patient's shame will not alter what a stranger believes. Both may be worth doing; neither substitutes for the other.

18.2 Public stigma

Public stigma is what a population believes, feels and does about people with mental illness. It is held by other people, it is distributed across a society and it is the type most people mean when they use the word without qualification. Its behavioural expression is **discrimination**, which is the part that has consequences: a tenancy refused, a job offer withdrawn, a marriage negotiation abandoned, a family member excluded from a decision.

Two distinctions matter for measurement, and both are frequently collapsed. The first is between an attitude and a behaviour: what someone reports believing about mental illness in the abstract predicts, but does not equal, what they do when a person with mental illness applies for a room. The second is between the general public and specific gatekeeping populations, notably employers, landlords, teachers, police and health professionals. Health professionals are a gatekeeping population in their own right, and stigma held inside the health system is the version with the most direct effect on whether a person gets treated at all.

18.3 Self-stigma

Self-stigma is what happens when a person who is aware of a public stereotype applies it to themselves and endorses it. The location has moved: this now lives inside the person, and its consequences are demoralisation, withdrawal from roles the person could still occupy, and reduced help-seeking. Its practical importance in a service is that it is a mechanism by which a

public phenomenon becomes a clinical one, and that it is largely invisible to any measure of public attitudes.

The service consequence is specific. Self-stigma degrades exactly the behaviours a chronic-illness service depends on: turning up, disclosing symptoms, accepting a rehabilitative role, applying for entitlements. A programme that reduces public stigma may leave self-stigma untouched, because a person can know that attitudes in general have improved and still hold the belief about themselves, so the two require separate measurement and separate intervention.

18.4 Structural stigma

Structural stigma is the version that requires no person to hold any attitude at all. It is the accumulation of rules, budgets, laws, institutional practices and physical arrangements that disadvantage people with mental illness, and it persists after every individual in the system has changed their mind. Recognising it is what turns a clinician into a public-health practitioner, because it is the only type of stigma whose remedy is administrative rather than educational.

Its examples in any health system are structural rather than attitudinal: mental health services funded at a lower share than the burden they carry, insurance products that exclude or limit psychiatric cover, facilities sited away from general hospitals so that attending one is itself a disclosure, records systems that make a psychiatric contact more visible than a medical one, and application forms that ask about psychiatric history in terms no medical condition attracts. Each of those is fixable by a decision, and none is fixable by a campaign.

TRAP **Answering a structural stigma question with an attitude intervention.** Asked how to reduce stigma affecting access to care, most candidates describe awareness work. Awareness addresses public stigma. Where the barrier is a funding rule, a siting decision or an insurance exclusion, the intervention is a change to that rule, and naming the rule is what distinguishes a public-health answer from a well-meaning one.

18.5 Associative stigma

Associative stigma, also called courtesy stigma, is the discredit that attaches to people connected to the affected person rather than to the person themselves. It falls on family members, on carers, and in some settings on the professionals who work in the field. It is the type most under-measured and, in a family-based care system, arguably the most consequential, because it operates on the people whom services rely on to deliver most of the care.

Its behavioural signature is concealment at the family level rather than at the individual level, with predictable effects on presentation delay, on treatment continuity and on marriage and employment decisions taken about other members of the household. A service that asks only the patient about stigma will not detect it.

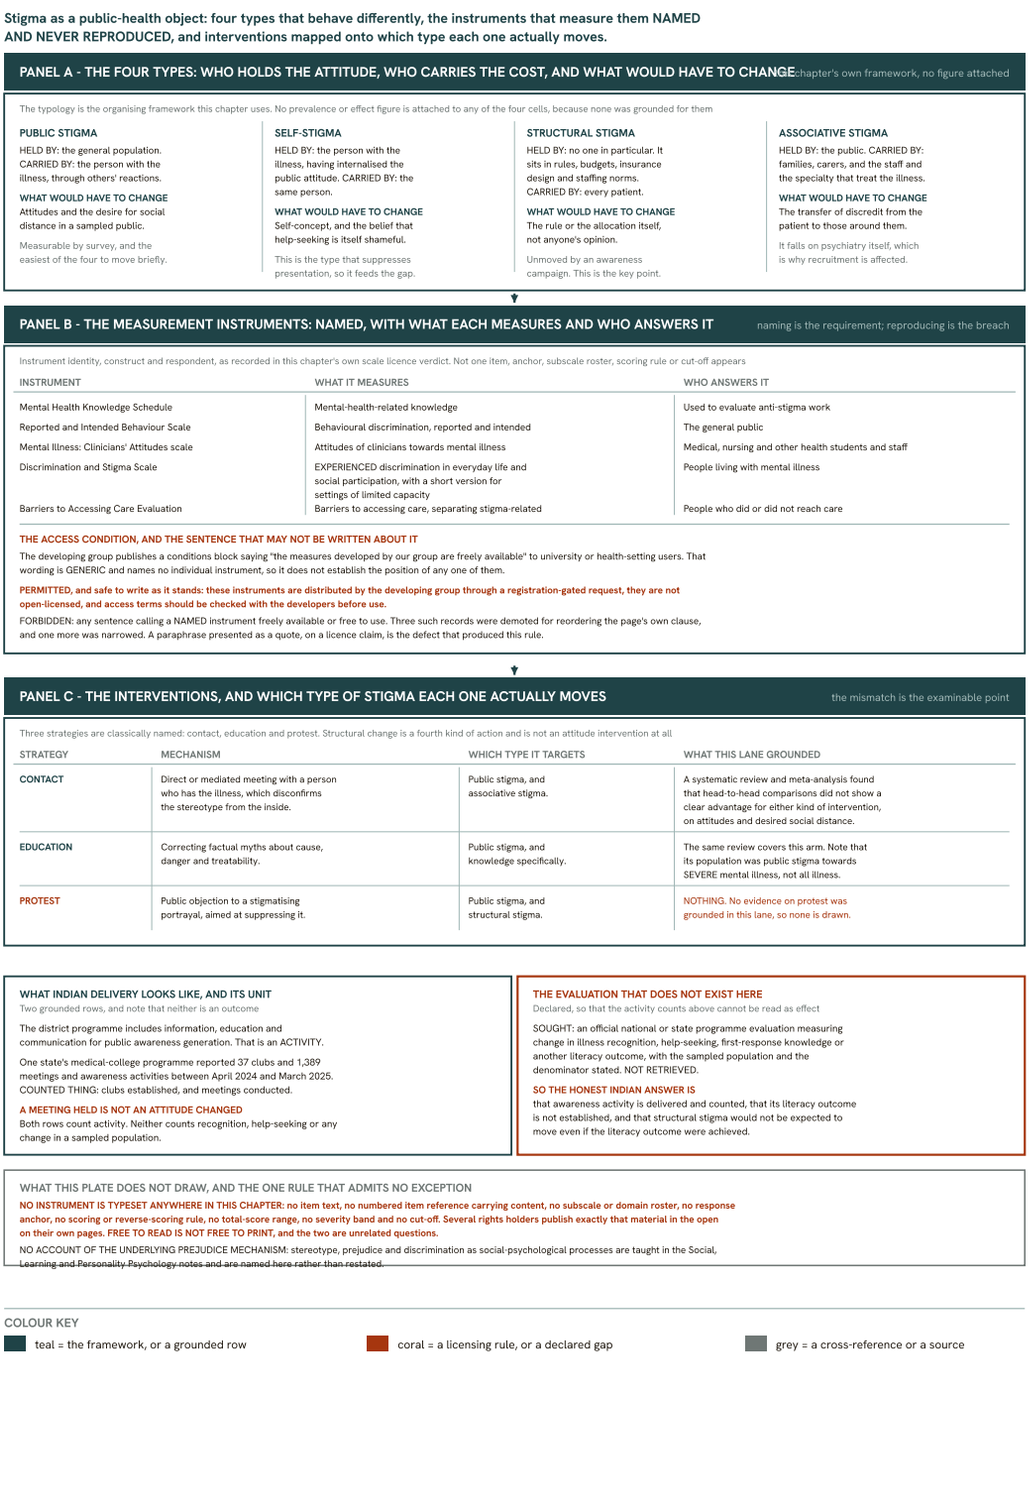


Fig 39.7 Stigma typed, measured and intervened on, with the measurement instruments named and none reproduced. The four types differ in who holds the attitude and who carries the cost, and structural stigma is the one an awareness campaign would not be expected to move. The instruments are named with their construct and respondent only. On access, the developing group's conditions block says its measures are freely available to university or health-setting users, but the wording is generic and names no individual instrument, so no named instrument may be called freely available; records that did so were demoted for reordering the page's own clause. On interventions, head-to-head comparisons showed no clear advantage for contact over education, and nothing on protest was grounded. Indian delivery is counted as activity, awareness generation through the district programme and 37 clubs with 1,389 meetings in one state between April 2024 and March 2025, while no official evaluation of any literacy outcome was retrieved.

MNEMONIC

Them, Me, The rules, Mine

The four locations of stigma, in the order in which a service should look for them. What others hold, what the person has taken on, what the system enforces without holding any view at all, and what falls on the family. Each location has its own intervention and its own measure, and an answer that covers all four has typed the problem rather than deplored it.

18.6 Where the clinical cut fits

A clinician assessing one patient will more often use a different cut of the same material, dividing stigma into what the person has actually met, what they anticipate meeting, and what they have internalised. That reading is taught in the Consultation-liaison, geriatric and transcultural psychiatry notes and is not re-derived here. What is worth stating in a public-health section is how the two cuts relate: the clinical cut is organised by the patient's relationship to the experience and is therefore useful in a consultation, while the population cut is organised by where the phenomenon lives and is therefore useful when choosing an intervention and an evaluation endpoint. They are not rival taxonomies and a candidate should not be trapped into defending one against the other.

18.7 Measuring stigma: what the instruments are for

Item by item this is the part of the topic candidates find hardest, because the instrument names are numerous and superficially interchangeable. They are not. Each was built to measure a specific construct, in a specific respondent, for a specific purpose, and using one in place of another produces a number that answers a question nobody asked. This section names them and describes them; it prints no item, no domain roster, no response option, no scoring rule and no threshold from any of them, and neither should you.

INSTRUMENT	CONSTRUCT IT MEASURES	WHO COMPLETES IT	WHAT IT CANNOT DO
Mental Health Knowledge Schedule (MAKS)	Mental health related knowledge in the general population	Members of the public	Detect discrimination, since knowing something is not doing something
Reported and Intended Behaviour Scale (RIBS)	Reported past and intended future contact behaviour toward people with mental illness	Members of the public	Establish that intended behaviour was enacted
Discrimination and Stigma Scale (DISC), and its ultra-short form (DISCUS)	Experienced discrimination across areas of everyday life and social participation	People with lived experience of mental illness	Measure what the public believes, since it records the receiving end only
Mental Illness: Clinicians' Attitudes scale (MICA)	Attitudes toward mental illness and psychiatry within a health workforce	Health students and health professionals	Stand in for public attitudes, since the respondent group is a gatekeeping profession
Internalized Stigma of Mental Illness scale (ISMI)	Internalised stigma in the affected person	People with lived experience of mental illness	Say anything about the surrounding population
Attribution Questionnaire (AQ) and Self-Stigma of Mental Illness Scale (SSMIS)	Attributional reactions to mental illness, and the self-directed form of the same process	Members of the public, and affected individuals respectively	Substitute for one another, despite sharing a theoretical origin
Community Attitudes toward Mental Illness (CAMI) and the Perceived Devaluation-Discrimination scale (PDD)	Community attitudes, and perceived societal devaluation as seen by the affected person	Community samples, and affected individuals respectively	Be read as measures of the same thing, since one records what a community holds and the other what a person believes a community holds
Barriers to Accessing Care Evaluation (BACE)	Barriers to accessing mental health care, distinguishing stigma-related barriers from other kinds	People who have needed or sought care	Attribute a barrier to a specific policy or institution

Two structural points about that table are worth more than the names in it. First, the respondent column is doing most of the work: some of these instruments interrogate the population and some interrogate the person receiving the population's behaviour, and a study that mixes them has measured two different phenomena and averaged them. Second, the last column is the one to reproduce in an answer, because knowing what an instrument cannot do is the evidence that you know what it does. Social distance measures, a family of instruments in this literature rather than a single named tool, sit alongside these and are usually treated as a behavioural proxy for public attitude.

GOOD TO REMEMBER

If you are going to evaluate anything you do about stigma, choose the respondent before you choose the instrument. If your intervention acts on staff, measure staff. If it acts on the public, measure the public. If your claim is that access improved, then the only respondent who can confirm it is the person who did or did not access care, and the instrument has to be one that asks them about barriers rather than one that asks anybody about attitudes.

18.8 Access to the instruments, and what is actually established about it

This is where a chapter can quietly go wrong, and it is worth being exact. The research group that developed a large part of the stigma instrument suite publishes a conditions block on each instrument's page. That block states that **the measures developed by the group are freely available to users who work or study in a university or health setting**, and it also forbids modification, forbids passing the measures to a third party and forbids charging for them. Every sentence of that block speaks of "the measures" generically and names no individual instrument.

The consequence is precise and it is the teaching point. Because the permissive sentence names no instrument, it does not establish the access terms of any particular one. Whether the clinicians' attitudes scales in that suite, the reported and intended behaviour measure, or the discrimination scale and its short form are individually free to a university or health-setting user is *not established* by the text that was sought and read. The defensible formulation, and the one to use, is that these instruments are distributed by the developing group through a registration-gated request, are not open-licensed, and that access terms should be confirmed with the developers before use. Other instruments in the field sit with individual rights holders or with commercial publishers and carry their own, different terms.

PITFALL

Assuming that a scale you can read online is a scale you may reproduce. Several rights holders publish item text, response options and scoring information openly on their own pages, for their own good reasons. Free to read is not free to print, not free to translate, not free to embed in a records system and not free to modify. A dissertation, a poster or a teaching handout that reproduces an instrument without checking its terms is a copyright problem regardless of how easy the material was to find.

18.9 Choosing an endpoint, which is the whole examination question

The last step is to connect type, instrument and endpoint into one line of reasoning, because that is what a viva question about stigma is really testing. If the target is public stigma, the endpoint is population knowledge, attitude or reported behaviour, and the measure is a public-facing instrument. If the target is the health workforce, the endpoint is professional attitude and the respondent is staff. If the target is self-stigma, the endpoint is internalised stigma in affected people. If the target is structural stigma, no attitude instrument will do at all, and the endpoint is the rule, the budget line or the entitlement itself, measured as a document rather than as a questionnaire. And if the claim is that any of this improved access, the endpoint is a barrier measure completed by people who needed care.

Anti-stigma interventions and their comparative evidence are taken up in their own section of this chapter, and mental health literacy, which overlaps with public knowledge but is a distinct programme class with its own outcomes, in the section after it. Neither of those names an instrument; the measurement layer stays here.

Location	Respondent	Rules	Not established
WHAT TYPES STIGMA: PUBLIC, SELF, STRUCTURAL, ASSOCIATIVE	WHAT DECIDES WHICH INSTRUMENT IS VALID	THE ONLY EVIDENCE THAT STRUCTURAL STIGMA MOVED	PER-INSTRUMENT ACCESS TERMS IN THIS SUITE

Type the stigma by where it lives, choose the respondent before the instrument, and remember that structural stigma is the only type whose remedy is a change to a rule and whose evidence is a document rather than a questionnaire.

19 · Anti-stigma interventions: contact, education and protest strategies and evidence

Anti-stigma work is a programme class, not a good intention. Three strategies are conventionally named, contact, education and protest, and a fourth family of structural measures sits alongside them without usually being counted as a strategy at all. Each has a design, a cost, a delivery workforce, a plausible mechanism, a characteristic failure mode and an outcome that can in principle be measured. This section treats them in exactly those terms, because a resident who can say what an anti-stigma programme would look like on the ground and how it would be evaluated has answered a public-health question, while a resident who can only assert that stigma should be reduced has not.

The social-psychological mechanism behind the contact strategy, the conditions under which intergroup contact reduces prejudice and the meta-analytic evidence establishing it, belongs to the Social, Learning and Personality Psychology notes and is named here rather than rebuilt. The stigma typology on which everything below depends, and the measurement instruments used to evaluate any of it, are taught in the section of this chapter that owns them; no instrument is named here.

19.1 Four families, and what each is actually trying to change

Before comparing the strategies it is worth being explicit that they act on different things, because most of the confusion in this topic comes from treating them as interchangeable ways of doing the same job.

- **Contact** puts members of a target audience into structured interaction with people who have lived experience of mental illness. Its target is the audience's beliefs and, more importantly, its behavioural intentions.
- **Education** supplies correct information in place of a false belief. Its target is knowledge, and its assumption is that stigmatising belief is partly an information deficit.
- **Protest** publicly objects to a stigmatising representation or practice and demands its withdrawal. Its target is not the audience at all, but the producer of the representation.

- **Structural measures** change a rule, a budget, a siting decision or an entitlement. Their target is the disadvantage itself, and they require no one to change their mind.

Two things follow immediately. First, only the last of the four addresses structural stigma, so a programme built entirely from the first three has left the most durable form of the problem untouched. Second, protest is aimed at a different recipient from the other three, which is why its evaluation looks different and why it carries a risk the others do not.

RULE **Match the strategy to the type of stigma and to the audience that holds it.** Contact and education act on people who hold attitudes. Protest acts on people who produce representations. Structural measures act on institutions. An intervention aimed at the wrong recipient can be delivered perfectly and still be incapable of changing the outcome it claims.

19.2 Contact as a programme rather than as a mechanism

Taking the mechanism as read, the programme questions are the ones a service actually has to answer. Who are the people with lived experience who will deliver it, how are they recruited, prepared and supported, and what happens to them afterwards? Contact-based programmes run on the disclosure of individuals, which is a genuine cost borne by the very people the programme exists to help, and a programme design that does not address the safeguarding, remuneration and sustainability of its speakers is exploiting them.

The other programme decisions are the audience, the dose and the medium. Audience matters because a gatekeeping group such as health staff, employers, landlords, teachers or police will produce a larger change in real-world discrimination per person reached than the general public will. **Dose**, meaning the number and spacing of encounters, matters because a single session is the commonest design and the least likely to produce durable change. Medium matters because live and mediated delivery differ in reach and in cost in opposite directions, and a filmed or written encounter can reach a population that no live programme could, at the price of some of what makes the encounter work.

19.3 Education as a programme

Educational anti-stigma work replaces a false belief with a correct one, and its natural home is any setting with a captive audience and an existing teaching structure: schools, colleges, professional training and workplace induction. Its advantages are that it is cheap, scalable, deliverable by people who are already employed to teach, and easy to standardise. Its limitations are equally clear. Knowledge and attitude are different constructs and correcting the first does not reliably move the second; attitude and behaviour are different again; and a message pitched to correct one stereotype can, if it is framed carelessly, reinforce the belief that the condition is a fixed and alien state.

The framing question inside educational work is genuinely contested and is worth knowing. Messages that present mental illness as an illness like any other are intended to reduce blame, and they can succeed at that while making the audience perceive greater difference and greater unpredictability. Messages that emphasise continuity between ordinary distress and disorder reduce perceived difference but can dilute the case for treatment. There is no framing that

optimises everything, and an examiner will be satisfied by a candidate who can name the trade-off rather than one who asserts a winner.

19.4 Protest, and why it is the least taught of the three

Protest challenges a specific stigmatising representation: an advertisement, a film, a headline, a product name, a public statement, an institutional practice. It is the only strategy of the three whose outcome is immediate, visible and binary, because either the material is withdrawn or it is not.

Its limitation is that suppressing a representation is not the same as changing the belief behind it, and there is a well-recognised risk of **reactance**, a defensive or contrary response in an audience that reads a public demand for withdrawal as censorship. It is therefore usually described as effective at the level of the representation and unreliable at the level of the attitude. In practice it is best understood as a tool for removing amplifiers rather than for persuading anyone, and it sits most comfortably alongside the structural family, because both act on institutions.

TRAP **Claiming attitude change from a protest outcome.** A withdrawn advertisement is evidence that a producer responded, not that a population's beliefs moved. The two are frequently reported together in programme write-ups, and separating them is exactly the discrimination an examiner is looking for.

19.5 Structural measures, which are where the durable gains are

Structural measures in anti-stigma work act on the rules. Removing an exclusion of mental illness from an insurance product, equalising an entitlement, siting a psychiatric service inside a general hospital, removing a psychiatric history question from a form that asks nothing comparable about physical illness, and correcting a funding allocation that is out of proportion to burden are all anti-stigma interventions, and none of them is a campaign.

Three features make them attractive and one makes them hard. They are durable, because a rule stays changed after enthusiasm fades; they are universal, because they apply to everyone the rule touches rather than to those who attended; and their evidence is a document, so evaluation is cheap and unambiguous. What makes them hard is that they require access to decision-makers rather than to audiences, which is a different competence and a longer timescale, and it is the competence a psychiatrist in a public system is uniquely placed to supply.

19.6 What the comparative evidence actually says

The question candidates most want answered is whether contact beats education. The best available synthesis is more useful and less satisfying than a winner: in a systematic review and meta-analysis of interventions to reduce stigma towards people with severe mental illness, **head-to-head comparisons did not show a clear advantage for either kind of intervention**. The outcomes in view were stigmatising attitudes and the desire for social distance, and the population in view was public stigma towards people with severe mental illness, defined in that review as schizophrenia, psychosis or bipolar disorder.

Read the scope before you use the conclusion. It concerns public stigma towards severe mental illness, not stigma towards common mental disorders, not self-stigma, not structural stigma and not the health workforce. It concerns short-term attitudinal and social-distance outcomes, not discrimination as experienced by patients and not access to care. A finding of no clear head-to-head advantage within that scope is not a finding that the two strategies are equivalent in every setting, and it is certainly not a finding that either changes behaviour.

The practical implication is liberating rather than deflating. If neither strategy is clearly superior on attitudinal outcomes, then the choice between them should be made on the grounds a programme manager can actually control: cost per person reached, availability of a willing and supported speaker group, existing delivery structures, the ability to reach the audience whose behaviour matters most, and the feasibility of measuring anything afterwards.

PITFALL

Reading "no clear advantage" as "both work well". A head-to-head comparison that fails to separate two interventions says nothing on its own about whether either produced a meaningful change against no intervention, and nothing at all about whether any change persisted or reached behaviour. Quote the comparison for what it is, a comparison, and keep the separate questions of efficacy, durability and behavioural transfer separate.

19.7 Targeting: everybody, or the people who decide

The last design decision, and the one with the largest effect on what a fixed budget buys, is whom to aim at. A whole-population campaign has the appeal of universality and the disadvantage that most of the people it reaches will never be in a position to discriminate against anyone with mental illness. A targeted programme aimed at a gatekeeping group accepts a smaller reach in exchange for a much higher density of consequential decisions per person reached, because every member of that group is repeatedly in a position to grant or withhold something: a job, a tenancy, a school place, an arrest decision, a referral, a bed.

Health professionals deserve a separate mention, because they are the gatekeeping group a psychiatric service can reach at zero marginal cost and the one whose stigma has the most immediate effect on whether a person is treated. Diagnostic overshadowing, in which a physical complaint in a person with known mental illness is attributed to the illness and not investigated, is the concrete form the problem takes inside a hospital, and it is a discrimination outcome rather than an attitude, which means it can be audited from records rather than surveyed. An anti-stigma programme aimed at one's own institution with an auditable endpoint of that kind is more evaluable than most of what appears in the published literature.

19.8 Endpoints, decay and the transfer problem

Three measurement problems recur across every anti-stigma programme and a good answer names all three.

The first is **endpoint selection**. Knowledge, attitude, intended behaviour, reported behaviour, experienced discrimination and access to care form a ladder, and each rung is harder to move and

more meaningful than the one below. Almost all published anti-stigma evaluation sits on the lower rungs, and the respondent has to change as you climb: the top of the ladder can only be reported by people who needed care.

The second is **decay**. Attitudinal change after a brief intervention tends to fade, so a design with a single post-intervention measurement taken immediately afterwards cannot distinguish a durable effect from a transient one. A follow-up measurement at a meaningful interval is not a refinement, it is the difference between an evaluation and a satisfaction survey.

The third is **transfer**. Attitude measured in a classroom does not automatically become behaviour in a rental market or a recruitment decision, and the gap between the two is where anti-stigma programmes most often fail without knowing it. Where a programme's real claim is about behaviour, the design has to include a behavioural or access-based endpoint, and if it cannot, the claim has to be scaled back to what was actually measured.

STRATEGY	STIGMA TYPE IT CAN MOVE	WHO RECEIVES IT	ENDPOINT THAT WOULD SHOW IT WORKED	CHARACTERISTIC FAILURE
Contact	Public, and workforce attitudes	An audience, ideally a gatekeeping one	Attitude and intended behaviour in that audience, remeasured after an interval	Single-session dose, unsupported speakers, and no follow-up measurement
Education	Public knowledge, and workforce knowledge	A captive audience in an existing teaching structure	Knowledge in that audience, and separately any attitude change	Knowledge improves while attitude does not, and the two are reported as one
Protest	The representation, not the belief	The producer of a stigmatising representation	Withdrawal or amendment of the material	Reactance in the audience, and attitude change claimed from a withdrawal
Structural measures	Structural stigma	An institution or a policymaker	The amended rule, budget line or entitlement, read as a document	Never attempted, because it needs access to decision-makers rather than to audiences

GOOD TO REMEMBER

If you are asked to reduce stigma in your own hospital, start with the workforce and with the rules, because those are the two things you can actually change and both act directly on whether a patient is treated well. Contact and education aimed at your own staff are deliverable inside existing teaching time, and the structural audit costs nothing: walk the pathway a psychiatric patient takes through your institution and write down every point at which it differs from the pathway a medical patient takes for no clinical reason. That list is an anti-stigma intervention plan.

MNEMONIC

Meet, Teach, Object, Amend

The four families in the order of increasing durability. Meeting someone changes an audience, teaching changes what an audience knows, objecting removes a representation, and amending a rule changes the disadvantage itself and stays changed. Most programmes stop at the first two, and most of the durable gains are in the last.

Audience	Producer	Institution	No clear advantage
WHAT CONTACT AND EDUCATION ACT ON	WHAT PROTEST ACTS ON	WHAT A STRUCTURAL MEASURE ACTS ON	CONTACT AGAINST EDUCATION, HEAD TO HEAD, PUBLIC STIGMA IN SEVERE MENTAL ILLNESS

Contact and education have not been separated head to head on public stigma towards severe mental illness, so choose between them on cost, audience and feasibility, and remember that the only strategy that reliably outlives the enthusiasm of the people running it is the one that changes a rule.

20 · Mental health literacy and community awareness programmes

A treatment gap is closed one step at a time, and the first step is inside somebody's head.

Before a person reaches a service they have to notice that something is wrong, decide that it is the kind of wrong a health system deals with, believe that something can be done about it, and know where to go. **Mental health literacy** is the name for that stock of knowledge and belief in a population, and awareness programmes are the deliberate attempt to raise it. This is a programme class in its own right, with its own audiences, its own delivery vehicles and its own outcomes, and it is not a synonym for anti-stigma work.

The distinction matters because the two are constantly conflated. Anti-stigma work targets attitude and discrimination and is taken up in its own section of this chapter. Literacy work targets recognition, help-seeking and first-response knowledge. They overlap, because a population that believes mental illness is a moral failing will not seek help for it, but they are separable, and a programme that measures one while claiming the other has produced an evaluation that answers the wrong question. No measurement instrument is named in this section; the instrument layer belongs to the section of this chapter that owns stigma measurement.

20.1 What literacy has to contain, derived from the pathway rather than asserted

Rather than reciting a list, it is more useful to derive the components from the pathway a person actually travels, because then each component has an owner and an intervention attached to it.

- **Recognition.** The person, or somebody near them, notices a change and names it as a change. Without this step nothing else happens, and this is the step at which a condition that presents as bodily complaint or as behaviour rather than as reported mood is most often missed.

- **Attribution.** The change is attributed to a health problem rather than to character, misfortune, external agency or the ordinary difficulty of life. This is the step at which explanatory models operate, and it is the step most sensitive to what a community believes about causes.
- **Efficacy belief.** The person believes something can be done. A population that recognises depression and believes it is untreatable will not convert recognition into help-seeking.
- **Route knowledge.** The person knows a specific place to go and what will happen there. "Seek help" is not route knowledge; a named service with a location, a cost and a plausible account of what the first visit involves is.
- **First-response knowledge.** The people around the person know what to do and what not to do in the interval before care is reached, which in this field includes knowing how to respond to talk of self-harm.

Reading the components this way has a practical payoff. It shows that a campaign can raise recognition and change nothing, because the bottleneck was efficacy belief or route knowledge, and it tells a programme designer to find out which step their population is actually stuck at before choosing a message.

RULE **Literacy is only useful up to the capacity of the service it points at.** A successful awareness campaign generates demand. If the district service cannot absorb it, the campaign converts an untreated population into a disappointed one, and the second is harder to reach the next time. Awareness and service capacity are commissioned together, or the awareness work should wait.

20.2 The Indian delivery vehicle that already exists

India does not need to invent a channel for this work, because the district programme already carries one. **Information, education and communication for public awareness generation is a named component of the District Mental Health Programme**, sitting alongside its service, training and other activities. That single fact is worth more in an examination than any number, because it locates the responsibility: public awareness is a programme deliverable with an owner in every district where the programme runs, not an optional add-on that depends on an enthusiastic individual.

Two consequences follow for a resident planning district work. The first is that the mandate already exists, so the question is not whether to do awareness work but how to do it well and how to show it was done. The second is that a component defined as information, education and communication invites the cheapest possible interpretation, which is the production of material. Producing a poster is an activity; changing what a population knows is an outcome; and the distance between the two is the whole difficulty of this topic.

IN INDIA

Because public awareness is a district programme component with a named owner, the decision a resident actually faces is what to do with a mandate that comes with no specified outcome. The defensible course is to choose one bottleneck in the pathway and one audience that can be reached repeatedly through an existing institution, then write down in advance what will be counted and what would count as success. A district that does this produces a small evaluable programme. A district that does not produces material, reports the material, and cannot say at the end of the year whether anybody knows anything they did not know before.

20.3 A worked state example, and what its numbers actually measure

State programmes supply the more instructive examples, because they report what they did. One state programme built around mental health clubs in government medical colleges reported, for a single year of operation, that **37 clubs were established across government medical colleges, conducting 1,389 meetings and awareness activities, between April 2024 and March 2025**.

Now read those two numbers with the discipline this chapter demands. The first counts *clubs established*, which is an institutional structure, and it is a real and useful thing to count because a standing structure inside an institution is what makes a programme survive the departure of the person who started it. The second counts *meetings and awareness activities conducted*, which is an output. Neither counts people, so neither is a reach figure; neither counts knowledge, so neither is a literacy outcome; and neither carries a denominator, so no proportion can be computed from either. What they establish is that a delivery structure existed and was active, which is a necessary condition for an effect and not evidence of one.

The generalisable lesson is that the institution is a good unit for this work. A college, a school, a workplace and a police training centre are all standing institutions with a captive audience, a teaching structure, a calendar and somebody accountable, which is why programmes built on institutions persist and campaigns built on events do not.

37	1,389	Output	Not established
CLUBS ESTABLISHED ACROSS ONE STATE'S GOVERNMENT MEDICAL COLLEGES	MEETINGS AND ACTIVITIES CONDUCTED IN THOSE COLLEGES IN ONE YEAR	WHAT BOTH FIGURES ARE, RATHER THAN REACH OR KNOWLEDGE	CHANGE IN RECOGNITION OR HELP-SEEKING, NATIONALLY OR IN ANY STATE

20.4 Designing a campaign that could work

Five design decisions determine whether an awareness intervention has any chance, and they should be made in this order.

Audience before message. A campaign aimed at everybody is aimed at nobody. The audiences worth choosing are those who are either at high risk, or in repeated contact with people at risk, or in a position to open the door to care. In practice that means adolescents and young adults, frontline health workers, teachers, employers and the police, and each requires a different message.

Bottleneck before content. Establish which pathway step the audience is stuck at. If recognition is intact and route knowledge is absent, a campaign explaining what depression is will change nothing that a card with an address would not have changed better.

Channel by audience, not by fashion. The channel that reaches a college cohort is not the channel that reaches an agricultural community or an older population, and the cheapest channel per contact is rarely the most effective per contact. Where a helpline or a service is named in the message, the channel and the service must be aligned so that the door is open when the message lands.

Language and idiom. A message has to arrive in the language people actually use for distress. Translation of a technical message into a regional language is necessary and not sufficient, because the idiom in which distress is expressed differs from the clinical vocabulary in every language. The cultural material behind this belongs elsewhere in this chapter and in the Consultation-liaison, geriatric and transcultural psychiatry notes; the programme point is simply that a message that does not use the population's own words for suffering will not be recognised as being about them.

Dose and repetition. A single exposure decays. Programmes embedded in a recurring institutional calendar accumulate; single events do not, however well attended they were.

A sixth consideration cuts across all five and is specific to this field. Awareness messages about mental illness can carry unintended harm in a way that messages about, say, hand hygiene cannot: material about suicide can affect a vulnerable reader, material that presents a condition as permanent can depress efficacy belief in exactly the people it hoped to reach, and material that names a diagnosis loosely can teach a population to apply a label rather than to seek an assessment. Every awareness message in this field therefore needs a review step that asks who the most vulnerable reader is and what this text does to them, and that step is cheap, is almost always omitted, and is the sort of thing a psychiatrist attached to a district programme is uniquely qualified to insist on.

PITFALL

Substituting an awareness day for an awareness programme. An annual observance is a communication opportunity, and it is often the only thing a service does. It reaches a self-selected audience once, it has no bottleneck analysis behind it, it produces an attendance figure rather than an outcome, and it recurs annually without accumulating. If a day is all the resource there is, attach it to a standing institutional structure so that it becomes the visible part of something continuous rather than the whole of it.

20.5 Measuring literacy, and what India does not know about itself

Literacy outcomes fall into an ordered set, and knowing the order is most of what an examiner wants. Activities delivered is the weakest, followed by people reached, then knowledge and beliefs in a sampled population, then help-seeking behaviour, then service contact, then outcomes among those who reached care. The rungs above knowledge require the person's behaviour rather

than their report, and the rungs above people reached require a sample and a denominator rather than a tally.

Here the position must be stated honestly. No official national or state programme evaluation measuring a change in illness recognition, help-seeking, first-response knowledge or any other literacy outcome, with the sampled population and denominator stated, was retrieved for India. That was sought through the national health mission, the national institute, and health-department sources for individual states. What is published is what programmes did, not what populations came to know. A resident asked what proportion of Indians can recognise depression, or whether awareness activity has changed help-seeking, should say that no such national figure is established, name what would be needed to establish it, and describe what is published instead.

TRAP **Reporting an activity count as evidence of impact.** Programme reports print what is countable, and what is countable is sessions, posters, participants and events. None of those is a literacy outcome. The correct sentence names the count for what it is, an output, and states separately that the outcome was not measured. Candidates who quote activity totals as achievements are repeating the error the report itself is careful not to make.

There is a second and quieter reason to insist on the ordered set. Programmes are funded on what they can report, so a funder that asks only for activity counts will get a system optimised for activity, and the people running it will not be behaving badly; they will be answering the question they were asked. Changing what is asked for is therefore a more powerful intervention on the quality of awareness work than any exhortation about rigour, and it is a change a clinician sitting on a district or state committee is in a position to propose. Ask for one population-sampled knowledge measure a year, with its denominator stated, and the programme will reorganise itself around producing one.

20.6 How this connects to everything else in the chapter

Literacy work sits upstream of almost every other topic here. It feeds the treatment gap, because recognition and route knowledge are two of the gap's determinants. It feeds suicide prevention, because first-response knowledge in the people around a person at risk is what a gatekeeper programme is trying to install in a workforce, and community awareness is the same intervention aimed at a population. It feeds primary care, because a population that presents earlier presents to a general clinic rather than to a specialist. And it is limited by service capacity, which is why it should never be planned in isolation from the service it points at.

PATHWAY STEP	WHAT THE POPULATION MUST HOLD	INSTITUTION THAT CAN SUPPLY IT	WHAT MEASURING IT REQUIRES
Recognition	The ability to name a change as a change	Schools, colleges, workplaces, frontline health workers	A sampled population and a stated denominator, not an attendance tally
Attribution	The belief that it is a health matter	Community channels in the population's own idiom	The same sample, asked about causes rather than about labels
Efficacy belief	The belief that something can be done	People with lived experience, and clinicians who are visible locally	Belief measured before and after, with an interval long enough to show decay
Route knowledge	A named place, a cost and an expectation of the first visit	The district programme's information component, and any helpline named in the message	Behaviour rather than report, ideally service contact attributable to the campaign
First response	What to do, and what not to do, before care is reached	Trained workforces and community groups around the person at risk	Knowledge in the responder, and where possible what the responder then did

Public awareness is a named component of the district programme, so the mandate exists everywhere the programme runs; what is published nationally is activity rather than knowledge, and no Indian literacy outcome with a stated population and denominator is established.

Psychiatric rehabilitation and community care

21 · Principles and components of psychiatric rehabilitation

Rehabilitation is what a service does about the part of illness that treatment does not reach. Symptoms can remit while a person remains unemployed, unhoused, unmarried, out of education and out of every role that gave their life shape, and no amount of additional medication addresses that. Psychiatric rehabilitation is the organised attempt to restore or build those roles, and read at population scale it is not an add-on to treatment but a distinct arm of the service with its own goals, its own workforce, its own outcomes and its own place in the architecture of a health system.

This section takes rehabilitation cross-diagnostically and at the level of systems: what the goals are, how disability and participation are framed and measured, who does the work, how the components are assembled into a network, and what has to happen for any of it to exist in a district. Rehabilitation as it is selected and delivered for an individual patient with schizophrenia, which is the disorder-specific psychosocial account, is taught in the Schizophrenia: management,

antipsychotics and adverse effects notes and is not restated here. Supported employment, deinstitutionalisation, the community settings and the Indian models each have their own section in this chapter and are named rather than taught here.

21.1 Three ideas that organise the whole field

Almost everything in rehabilitation follows from three constructs, and the reason candidates find the field vague is that they learn its activities without learning these.

- **Recovery** in this literature is not synonymous with remission. It denotes a person living a life they regard as worth living, with or without continuing symptoms, and it is defined by the person rather than by the clinician. That shift in who defines the endpoint is the single most consequential idea in modern rehabilitation, because it makes the goal negotiable rather than prescribed.
- **Disability** is the gap between what a person can do and what their environment demands and permits. It is therefore not a property of the person alone, and it can be reduced by changing the environment as readily as by changing the person, which is what makes it a public-health object rather than only a clinical one.
- **Participation** is the outcome that both of the above are aiming at, and it has a usable definition: participation means joining in community activities. That definition is deliberately mundane, and its mundanity is the point. It is checkable, it is meaningful to the person and their family, and it does not require an instrument to notice.

The international orientation of the field follows the same logic. **The World Health Organization's guidance on community mental health services sets out the case for shifting away from institutional, biomedical models of care and toward person-centred, recovery-oriented and human-rights-based approaches**. Each of those three adjectives corresponds to one of the constructs above, and together they describe not a technique but an orientation that a whole service either has or does not.

RULE **The rehabilitation goal is a role, not a score.** A plan whose stated aim is symptom reduction or an improved rating is a treatment plan wearing a rehabilitation label. A rehabilitation aim names the role to be resumed or acquired, the setting in which it happens, and the supports required to sustain it, and it is agreed with the person whose role it is.

21.2 The disability frame, and why psychiatry borrows it

Psychiatry does not need a private theory of disability, and using the general health one has three advantages: it is cross-diagnostic, it is shared with the rest of medicine and with the disability sector, and it is the frame in which entitlements are written. **The International Classification of Functioning, Disability and Health supplies that framework, and it is explicitly designed to apply at both the individual and the population level**, which is precisely what a chapter about services needs: the same construct that describes one person's limitation aggregates into a description of a district's burden.

Its companion assessment instrument is the **World Health Organization Disability Assessment Schedule, WHODAS 2.0**, and its defining property for our purposes is scope: it is used across all diseases, including mental, neurological and addictive disorders. That cross-diagnostic reach is why a rehabilitation service can use one functioning framework for a mixed caseload rather than a different measure for every diagnosis, and it is why functioning data from a psychiatric service can be compared with, and pooled with, functioning data from the rest of the health system.

The instrument is named here and not reproduced. Its domain structure, its scoring and its administration belong to the chapter that owns assessment instruments, the Psychotherapies, clinical assessment, MSE and nosology notes, and no part of it is printed in this chapter. What a resident needs from it here is the conceptual point: functioning is measurable, it is measurable in the same terms across conditions, and it is a different quantity from symptom severity.

21.3 Community inclusion as the outcome, and as an entitlement

Set against a history in which people with mental illness were removed from society for years at a time, the modern statement of the rehabilitation aim is unusually direct. In India it is not merely a professional aspiration but a legal one: **under the Mental Healthcare Act, 2017, every person with mental illness has the right to live in, be part of, and not be segregated from society**. The architecture of that statute, the machinery of rights, capacity, admission and review that surrounds the provision, belongs to the Mental health legislation and forensic psychiatry notes and is not rehearsed here.

What matters for a service is the direction of the obligation this creates. If community living is an entitlement, then segregation requires justification rather than inclusion requiring permission, and a person occupying an institutional bed for want of anywhere to go is not a clinical fact but a service failure with a legal character. That reframing is what converts rehabilitation from a benevolence into a duty, and it is the argument a resident should be able to make in one sentence.

IN INDIA

Two Indian facts sit together and they generate a concrete decision. Community living is a statutory entitlement, and the national programme's stated objectives include promoting community participation in the development of mental health services and stimulating efforts toward self-help in the community. Together they mean that when a resident writes a rehabilitation plan in a public service, the community component is not discretionary and is not something to be added if resources allow. The operational consequence is that the plan should name the community role, the person or group in the community who will support it, and the district programme staff member responsible, before it names any additional treatment.

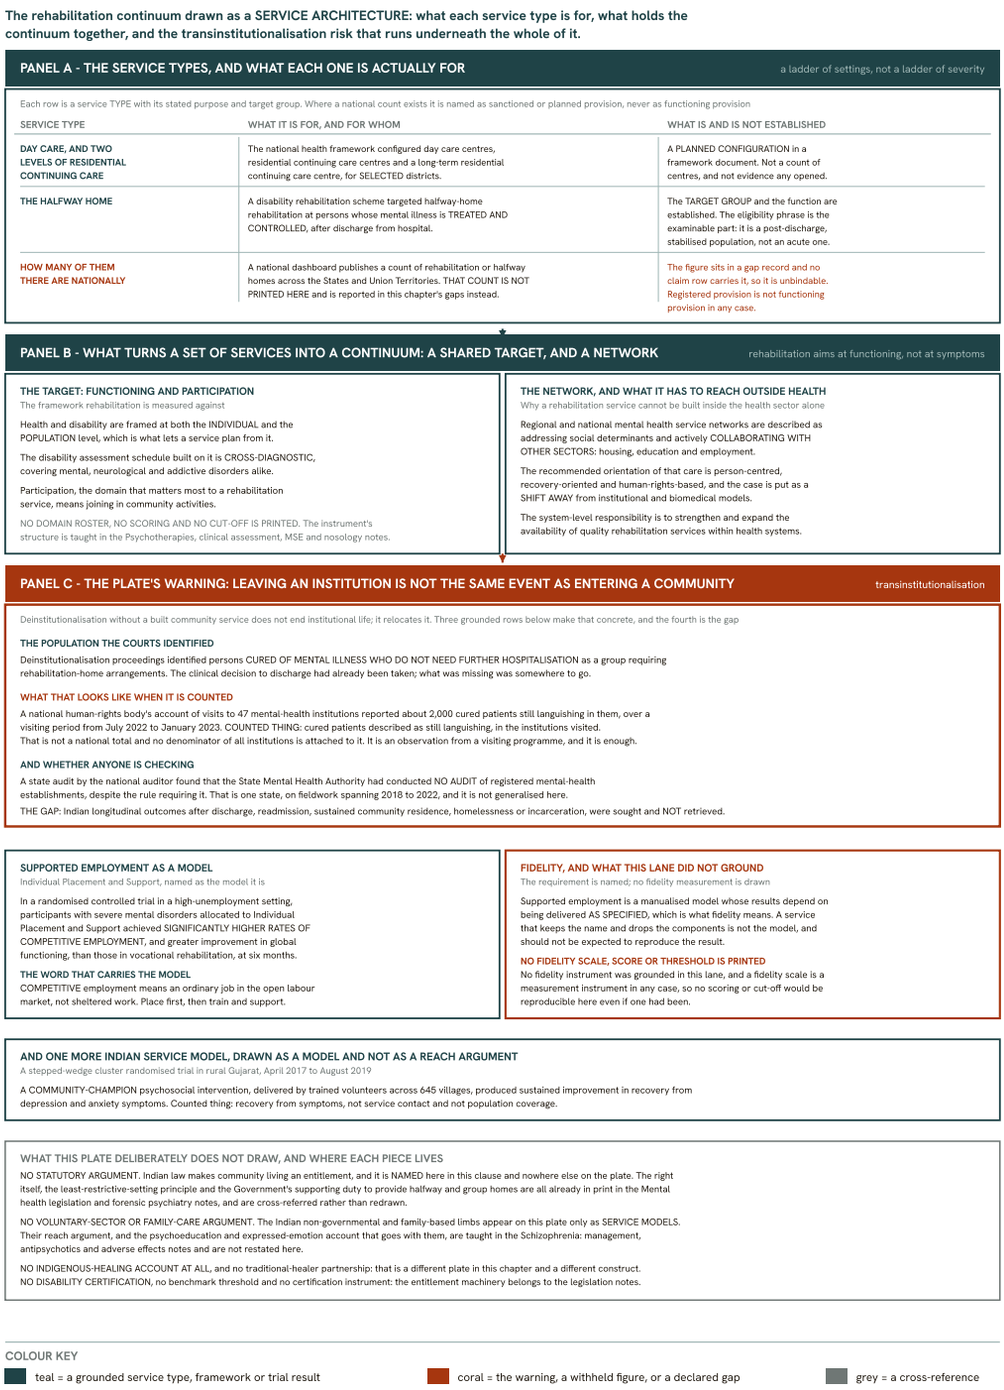


Fig 39.8 The rehabilitation continuum as a service architecture, with transinstitutionalisation as its warning. The national framework configured day care and two levels of residential continuing care for selected districts, and a disability scheme targeted halfway-home rehabilitation at people whose mental illness is treated and controlled after discharge; the national count of such homes sits in a gap record that no claim row carries, so it is not printed. The continuum is held together by a target, participation, meaning joining in community activities, and by service networks that collaborate across housing, education and employment. The warning is concrete: courts identified cured people who no longer needed hospitalisation but had nowhere to go, a human-rights body reported about 2,000 such patients still in the 47 institutions it visited, one state auditor found no audit of registered establishments had been done at all, and Indian post-discharge outcomes were sought and not retrieved. Supported employment is drawn as a model: it achieved significantly higher competitive-employment rates than vocational rehabilitation in a randomised trial, and its fidelity requirement is named without any fidelity instrument being reproduced.

21.4 The components: three ways to close the same gap

If disability is the gap between what a person can do and what the environment demands, then only three kinds of action can close it, and every rehabilitative component in the field is one of the three. This is the most useful classification a resident can carry, because it converts a shapeless list of activities into a decision.

Raise the person's capacity. Skills training in the domains the role requires, whether that is self-care, household management, money, travel, conversation, or the specific competencies of a job or a course. This is the component people think of first, it is the slowest of the three, and it is the only one that transfers with the person when they move.

Supply a support that substitutes for the missing capacity. A person, a structure or a piece of assistance that does what the person currently cannot: a family member who manages medication, a supported setting, a job coach, a case manager, a prompt, an entitlement that replaces lost income. Supports work immediately and stop working when they are withdrawn, which is why the plan has to state whether a support is a bridge or a permanent adjustment.

Lower what the environment demands. Modified hours, an adjusted course load, a changed task, a physical relocation of a service closer to where people live, a form simplified, a rule waived. This is the cheapest of the three and the one clinicians reach for last, because it requires negotiating with somebody outside the health system rather than doing something to a patient.

Most real plans use all three, and the useful discipline is to say which is which, because they have different timescales and different failure modes. A plan built only of skills training will show nothing for months; a plan built only of supports will collapse when the support is withdrawn; and a plan built only of environmental adjustment leaves the person no more able than before if the environment changes again.

21.5 Who does rehabilitation, and what a team has to contain

Rehabilitation is the least doctor-centred part of psychiatry, and the composition of the team is not incidental to it. In India the district programme is built around a multidisciplinary team: **the district team under the programme includes a psychiatrist, a clinical psychologist and a psychiatric social worker**. Those three disciplines map onto three different problems, and understanding the mapping is more useful than memorising the list. The psychiatrist holds diagnosis, medication and risk; the clinical psychologist holds assessment of functioning and psychological intervention; the psychiatric social worker holds the person's relationship with their family, their entitlements and their community, which is where most rehabilitation actually happens.

Two riders follow. First, a team roster is a statement about establishment and not about occupancy: a sanctioned post is not a person, and a district with vacancies in two of the three disciplines has a team on paper only. Second, the roster above is what the programme names and is not asserted here as the complete list of anybody's staffing; occupational therapy, nursing, community health workers, peer workers and vocational staff all have roles in a mature rehabilitation service, and the absence of any of them shapes what that service can offer.

TRAP **Listing rehabilitation activities in place of rehabilitation principles.** Asked about psychiatric rehabilitation, most candidates produce a list of interventions. An examiner is looking for the organising logic first: a person-defined goal expressed as a role, disability understood as the fit between person and environment, participation as the outcome, and a team whose composition follows from the fact that most of the work is social. The activity list is worth marks only once that frame is on the page.

21.6 From a team to a system: networks and other sectors

A team can rehabilitate a person only as far as its own walls. Everything a person needs to occupy a role, somewhere to live, something to do, money, transport and a social world, is held by other sectors, which is why rehabilitation is the part of psychiatry most dependent on institutions that do not report to health. This is not a rhetorical point but the accepted design: comprehensive mental health service networks are described as addressing the key social determinants of health and actively collaborating with other sectors such as housing, education and employment.

The practical form of that collaboration is unglamorous and specific: a named contact in the social welfare office, a working relationship with a disability certification process, a school or college willing to take a student back on modified terms, an employer prepared to make an adjustment, a landlord or a hostel that will accept a tenant. A rehabilitation service is largely a set of such relationships, and the ones that are written down and institutionalised outlive the individuals who built them.

21.7 Getting it to exist: implementation as the real problem

Rehabilitation suffers from a characteristic implementation deficit. It is uncontroversial, it appears in every policy document, and it is the first thing cut when a service is under pressure, because its absence produces no immediate crisis. The international guidance addresses exactly this: the recommendations on rehabilitation in health systems are addressed to member states and other stakeholders in order to strengthen and expand the availability of quality rehabilitation services, which is to say the problem being solved is availability rather than technique.

What has to be true for a district rehabilitation service to exist, in order of how often each is missing: a named person whose job it is; a physical space that is not a ward; a link to at least one non-health sector; a referral route in and a discharge route out; and a record of who is in the service and what they are working toward. Notice that none of those is a treatment, and that a service with all five and no new intervention will do more rehabilitation than a service with a manual and none of them.

21.8 The frameworks compared, and what each is good for

The commonest confusion in this topic is between the constructs, the frameworks and the instruments, so it is worth setting them side by side with what each can and cannot deliver.

WHAT IT IS	WHAT IT FIXES	LEVEL IT OPERATES AT	WHAT IT CANNOT DO
Recovery orientation	Who defines the goal, and therefore what counts as success	The individual, and the culture of a whole service	Be measured directly by any symptom or disability score
The functioning and disability classification	A shared vocabulary for limitation and for the environment's contribution to it	Individual and population, by design	Tell you what to do about any particular limitation
The cross-diagnostic disability assessment schedule	Comparable measurement of functioning across conditions	The individual, aggregable upward	Substitute for a symptom measure, or for the person's own account of their goals
The rights framework in statute	The direction of the obligation: inclusion is the default and segregation needs justifying	The service and the state	Create the community capacity that makes the entitlement usable
Service networks with other sectors	Access to housing, education, employment and income, which sit outside health	The system	Function without named relationships and someone to maintain them

GOOD TO REMEMBER

In an outpatient clinic the whole of this section reduces to one extra question at the end of a review: what does this person do all day, and what would they like to be doing that they currently cannot? A person whose symptoms are controlled and whose answer to that question has not changed in two years is not doing well, whatever the mental state examination records, and that is the moment to write a rehabilitation aim rather than to adjust a dose.

21.9 Evaluating rehabilitation, and the honest position

Rehabilitation is evaluated on roles occupied and sustained rather than on symptoms: living situation, occupation or education, income and entitlements received, social contact, and functioning measured on a cross-diagnostic scale. Two design cautions apply. Outcomes take time, so a study window shorter than the time it takes to acquire a role will find nothing regardless of the intervention. And the environment is part of the outcome, so a service operating in a labour market or a housing market that offers nothing will produce poor role outcomes with excellent practice, which is an argument for reporting the setting alongside the result rather than for despair.

MNEMONIC

Goal, Gap, Team, Sector, Place

The five things a rehabilitation plan has to name. The person's own goal expressed as a role; the gap between what they can do and what the environment demands; the discipline in the team who will hold the work; the sector outside health whose cooperation is needed; and the physical place where it happens. A plan missing any one of them will not survive contact with a district.

Role	Cross-diagnostic	Statutory	Availability
THE UNIT OF A REHABILITATION GOAL, NOT A SYMPTOM SCORE	SCOPE OF THE DISABILITY ASSESSMENT FRAMEWORK USED	STATUS OF COMMUNITY LIVING FOR PEOPLE WITH MENTAL ILLNESS IN INDIA	THE PROBLEM THE INTERNATIONAL GUIDANCE ADDRESSES, NOT TECHNIQUE

Rehabilitation aims at a role the person chooses, works on the gap between the person and their environment, measures functioning on a framework shared with the rest of medicine, and in India rests on a statutory entitlement to live in and not be segregated from society.

22 · Individual Placement and Support (IPS) and supported employment

Work is the rehabilitation outcome patients ask for and services most often decline to organise. Ask a person with a severe mental illness what they want and paid work appears near the top of the list far more often than clinicians expect. Work supplies income, structure, identity, social contact and a reason to be somewhere at a fixed time, all of which are exactly the things illness takes away, and it is the outcome most visible to the family. Supported employment is the service class built to deliver it, and Individual Placement and Support is the model that dominates the field.

This section treats the model as a programme: how a vocational service is designed, what its implementation barriers are, what its fidelity means, how it is evaluated, and what its Indian position actually is. The model as an intervention chosen and delivered for one patient with schizophrenia, and the older distinction between preparing a person before placing them and placing them and then supporting them, are taught in the Schizophrenia: management, antipsychotics and adverse effects notes and are named here rather than re-derived. The principle set and fidelity criteria published in the model's own manual are not reproduced in this chapter.

22.1 Why a health service is in the employment business at all

The objection is worth answering directly, because an examiner may raise it. Employment is not a health service's core competence, employers are not health institutions, and a labour market is not something a district hospital can influence. Three arguments answer it.

- **The barrier is clinical in origin.** The reason many people with severe mental illness are out of work is not a lack of jobs in general but a set of consequences of illness and its treatment,

together with discrimination. Nobody else is positioned to address those, and a generic employment service that does not understand them will fail with this population.

- **The support has to be coordinated with treatment.** Medication timing, appointment scheduling, relapse recognition and the handling of a period of illness while employed all require the clinical team, and a vocational service operating with no line to the clinicians will lose people at the first setback.
- **Work changes the clinical picture.** Income, structure and role are not incidental to outcome, which is why employment is a rehabilitation aim and not a life event that happens after treatment finishes.

RULE **The endpoint is competitive employment, not activity.** **Competitive employment** means an ordinary job in the ordinary labour market, open to anyone, paid at the going rate. Occupational activity within a service, sheltered production and unpaid work experience are different endpoints entirely, and a programme that reports them as employment has changed the outcome definition without saying so.

22.2 The two service logics, and what each demands of a system

Vocational rehabilitation has historically been organised around preparing a person first, through assessment, training and graded work in a protected setting, and placing them once they are judged ready. **Supported employment** inverts the order: the person is helped into an ordinary job early and the support is delivered on the job. The contrast itself is taught in the notes named above; what belongs here is what each logic requires a service to build, because that is what makes one of them feasible in a district and the other not.

A preparation-first system requires physical infrastructure, a sheltered setting, equipment, supervision and a production process, and its unit of investment is a building. A place-and-support system requires almost no infrastructure and a great deal of relationship: employers who will take a referral, a specialist who can spend the day travelling, and a clinical team willing to be interrupted. The first is easier to fund because it is visible and countable; the second is harder to fund and reaches ordinary employment.

The distributional consequence is the one worth stating in an answer. A preparation-first system has an implicit gate, since somebody decides who is ready, and gates exclude the people with the most severe illness. A place-and-support system with an open door reaches those people and accepts that some placements will fail. Which of those a service prefers is a value choice about who the service is for, not a technical one.

22.3 The design decisions that define a supported-employment service

Individual Placement and Support is the best-specified expression of the place-and-support logic. Rather than reciting its published principles, it is more useful to set out the decisions any vocational service has to make, because those decisions are what a resident will be asked about and what a service actually gets wrong.

DESIGN DECISION	THE RESTRICTIVE OPTION	THE SUPPORTED-EMPLOYMENT OPTION	WHAT THE CHOICE COSTS OR BUYS
Who is eligible	Referral after an assessment of readiness	Anyone who says they want to work	An open door reaches severe illness and raises the failure rate; a gate protects the success rate by excluding the people the service exists for
What counts as a job	Sheltered or in-service occupational activity	An ordinary paid job in the open labour market	Sheltered work is easier to arrange and does not produce the income, status or social contact that the outcome is valued for
When the search starts	After a period of preparation and training	Early, with support delivered on the job	Preparation delays and loses people at every step; early search accepts that learning happens in the real setting
Whose preference decides	What the service can arrange	The kind of work the person actually wants	Jobs chosen by the person are held longer; jobs chosen for them are convenient to arrange and quickly lost
Where the employment worker sits	In a separate vocational agency	Inside the clinical team, sharing its meetings and its record	Integration is the expensive part and is repeatedly the part dropped, after which the service is no longer the model it is named for
How long support lasts	Until placement, then discharge	Continuing as long as it is needed, on both sides of the placement	Time-limited support produces placements and not retention; continuing support is a standing cost with no dramatic moment to justify it
Whether income advice is included	Not the service's business	Part of the work, because entitlements interact with earnings	Without it a person can be advised into a job that leaves them worse off, and will reasonably refuse the next offer

Reading the table downward gives the model. Reading it across gives the examination answer, which is that each decision has a defensible restrictive version and that the supported-employment position is a coherent set of choices rather than a slogan.

22.4 The evidence, and how to read a single trial

The comparison that matters is against conventional vocational rehabilitation rather than against nothing. In a randomised controlled trial conducted in a high-unemployment setting, **participants receiving Individual Placement and Support achieved significantly higher rates of competitive employment and showed greater improvements in global functioning than participants receiving vocational rehabilitation**, with those outcomes assessed at six months in people with severe mental disorders who met the trial's eligibility criteria.

Four things in that sentence are load-bearing and a candidate should be able to point to each. The comparator is an active one, so the finding is not merely that doing something beats doing nothing. The primary outcome is competitive employment, which is the endpoint that matters rather than a proxy. A second outcome, global functioning, moved in the same direction, which is the finding that argues against the objection that employment gains come at a cost to the person's mental state. And the setting was one of high unemployment, which matters for the next subtopic.

TRAP **Quoting an employment rate as though it were a property of the model.** Employment rates achieved in vocational trials are functions of the model, of the eligibility criteria and of the labour market the trial ran in, and the last of the three moves the number more than the first. A rate quoted without its labour market is uninterpretable, which is why this section reports the direction of the comparison and prints no employment percentage at all.

22.5 Transportability, which is the real Indian question

Transportability, meaning whether a result obtained in one place will hold in another, is the question this evidence raises most sharply, because every employment outcome is defined against a labour market and the first thing to ask of any vocational evidence is what market it came from. A trial in a high-unemployment setting is informative in a specific way: it tests the model where jobs are scarce, which is a harder test than a tight labour market provides, and it therefore travels better than a result obtained where employers are competing for staff.

The Indian question is not only about the quantity of work but about its form. Where a large share of employment is informal, seasonal, self-organised or family-based, the definition of competitive employment, an ordinary job in the open market at the going rate, describes a smaller part of the working population than it does in the settings where the model was developed and tested. That does not make the model inapplicable; it means the outcome definition has to be examined before the model is imported, and that a service which counts only formal-sector placements will under-record what it achieved. Work within a household or a family enterprise is a real vocational outcome in this setting, and it is taught in the notes cross-referred above rather than here.

The honest position on Indian implementation is that this chapter's sources carry no coverage figure, no service count and no outcome data for supported employment in India. Formal vocational services organised on this model are scarce, and a resident asked about Indian supported employment should say that the model's evidence base is international, that no

national Indian implementation figure is established, and that the transportability question is about the labour market and the outcome definition rather than about the intervention.

22.6 Fidelity, and why the word matters

Fidelity is the degree to which a service as delivered matches the model as specified. It matters in this field more than in most, because supported employment is easy to name and hard to run, and the components that get dropped under pressure are systematically the expensive ones: the integration of the employment worker into the clinical team, the open door, and the continuation of support after placement. A service that has dropped those is still called by the model's name, and its results are then reported as results of the model.

Two consequences follow. For a reader of the literature, an implementation study that reports poor outcomes without reporting fidelity has not shown that the model failed. For a service, fidelity is the thing to measure when outcomes are disappointing, before concluding that the approach does not work locally. Structured fidelity instruments exist for this model and are published by the groups that developed them; none is reproduced here, and a service intending to use one should obtain it from its rights holder.

> **PITFALL**
>
> Calling a job placement scheme by the model's name because it places people in jobs. The name denotes a specified set of components, and the two that are almost always missing in practice are the employment worker's membership of the clinical team and the continuation of support after the person starts work. A service without those places people and then loses them, which is not the model performing badly but a different service performing as expected.

22.7 Evaluating a vocational service

The endpoints form an ordered set and each answers a different question. Number entering the service answers a reach question. Number placed in competitive employment answers an efficacy question. Time to first placement answers an efficiency question. **Job tenure** and total time worked answer the retention question, which is the one that distinguishes a service from an introduction agency. Earnings answer the question the person actually asked. And functioning and clinical outcome answer the objection that work will destabilise people.

Two cautions apply to all of them. The denominator has to be stated, since a placement rate computed over people who completed a preparation phase is not comparable with one computed over everybody who walked through the door. And the follow-up has to be long enough for retention to be visible, because a placement measured at the moment it happens tells you nothing about whether the person is still there.

There is one further evaluation trap peculiar to this field and it is worth naming. Employment outcomes are highly sensitive to who is counted and when the counting starts, so two services with identical practice can report very different results by defining their cohort differently: everyone referred, everyone who attended once, everyone who set a work goal, or everyone still engaged at the point of analysis. The last of these is the most flattering and the most common,

because it silently removes the people for whom the service did not work. A published vocational result should always be read backwards from the denominator to the referral, and a service reporting its own results should state the referral cohort before anything else.

GOOD TO REMEMBER

Ask about work at every review, in the same routine way as sleep and appetite, and ask what the person wants rather than what they are ready for. In a service with no vocational specialist the useful actions are small and available: agree a work goal in the notes, schedule appointments so they do not collide with a working day, avoid regimens that make morning work impossible where an alternative exists, and speak to the family about the difference between protecting somebody and keeping them idle.

MNEMONIC

Open door, Ordinary job, Own choice, One team, Ongoing support

Five things a supported-employment service either has or has not. No readiness gate, an ordinary paid job in the open market, the work the person chose, the employment worker inside the clinical team, and support that does not stop at placement. When outcomes disappoint, check these before questioning the approach.

Competitive	**Higher**	**Retention**	**Not established**
THE ONLY EMPLOYMENT ENDPOINT THAT COUNTS AS THE OUTCOME	EMPLOYMENT WITH THE MODEL AGAINST VOCATIONAL REHABILITATION, RANDOMISED COMPARISON	WHAT SEPARATES A SERVICE FROM AN INTRODUCTION AGENCY	ANY INDIAN COVERAGE OR OUTCOME FIGURE FOR THIS MODEL

Supported employment places the person in an ordinary paid job first and supports them there, its endpoint is competitive employment and its evidence is a randomised comparison against conventional vocational rehabilitation, and what fails in practice is almost always fidelity rather than the model.

23 · Deinstitutionalisation: rationale, outcomes and lessons

A service transition, not an event and not a policy announcement. **Deinstitutionalisation** names the deliberate shift of the centre of gravity of psychiatric care out of large custodial institutions and into general hospitals, community teams and ordinary housing. Every country that has attempted it has discovered the same thing: closing beds is easy, and building the community capacity that has to receive the people is slow, expensive and politically invisible. The distance between those two facts is where the entire teaching content of this topic lives.

The statutory right that mandates community living, and the rights architecture that surrounds it, belong to the Mental health legislation and forensic psychiatry notes and are named here rather than rehearsed. What this section owns is the transition itself: why it was undertaken, what its outcomes were, what has to exist for it to succeed, what happens when the second half is not

built, and what the Indian record actually documents. The community settings that receive people, the day care centre and the halfway home, are configured in their own section of this chapter.

23.1 What drove it, and why the drivers matter

Four forces converged, and it is worth separating them because they pull in different directions and because a service that understands only one of them will make characteristic mistakes.

- **Therapeutic.** Effective treatment shortened the period during which continuous institutional supervision was necessary, so a category of person appeared who no longer required a hospital and had nowhere else organised to be.
- **Ethical and legal.** Long-stay institutions were repeatedly found to produce harm of their own, and the growth of rights-based thinking made indefinite custodial care without justification indefensible rather than merely regrettable.
- **Clinical.** Prolonged institutional living produces its own disabilities, a loss of initiative, of skills and of social contact, which are additional to the illness and are caused by the setting.
- **Economic.** Large institutions are expensive to run, and this is the driver that most often becomes the operative one. That matters, because a transition motivated by cost has a strong incentive to complete the closure and a weak incentive to fund the replacement.

RULE **Deinstitutionalisation is a two-sided programme and only one side is self-executing.** Bed closure happens by administrative decision. Community capacity happens only if somebody funds, staffs and builds it. A programme that treats the second as a consequence of the first has not planned a transition, it has planned a discharge.

23.2 The arc, told as a pattern rather than as a chronology

The international experience repeats a recognisable sequence, and knowing the sequence is more useful than knowing any one country's dates. It begins with a long asylum era in which the institution is the service and admission is effectively indefinite. A critique then accumulates from several directions at once: from clinicians who observe that a portion of the disability they are treating is produced by the setting, from lawyers and rights advocates who object to indefinite detention without review, from the emergence of treatments that make continuous supervision unnecessary for many people, and from finance ministries that notice the running cost of large estates.

Bed reduction then proceeds, usually faster than the community programme that was supposed to accompany it, and the visible consequences appear in other sectors rather than in psychiatry: in homelessness, in criminal justice, in nursing and residential homes, and in the burden carried by families. A later corrective phase follows in which community teams, crisis services, supported housing and assertive follow-up are built, and the systems that build them find that community care is not obviously cheaper than the institutions it replaced, only better, which is a very different argument to make to a treasury.

The clinical construct that justified the whole enterprise deserves its own name. **Institutionalism** is the syndrome produced by prolonged institutional living itself, comprising loss of initiative and of ordinary decision-making, submissiveness to routine, loss of practical and social skills through disuse, and a narrowing of expectation until the person no longer wants what they cannot have. It is additional to the illness, it is caused by the setting, and it is partly reversible, which is why it is an argument for moving people rather than an argument that they cannot be moved.

23.3 What has to exist on the receiving side

The precondition list is short and each item is a place where real programmes have failed.

Somewhere to live. Housing is the precondition, not an adjunct, and it is the one a health department cannot supply on its own. A discharge plan without an address is a transfer of the person to their family or to nowhere.

An income. A person leaving a long stay has usually lost earning capacity, savings and often any claim on family property. Without an income or an entitlement, community living is not sustainable regardless of clinical state.

A team that follows. Continuity of clinical care after discharge, with somebody responsible for finding the person when they do not attend, is what distinguishes community care from discharge. The absence of an **assertive follow-up** function, meaning a service that goes after the person rather than waiting for them, is the commonest technical failure.

Something to do. Occupation and role, addressed elsewhere in this chapter, without which the community placement is a smaller and lonelier institution.

A family or a substitute for one. In a system where families provide most long-term care, the family's willingness and capacity is a service variable and not a private matter, and it has to be assessed rather than assumed.

23.4 Transinstitutionalisation and the other failure modes

Transinstitutionalisation is the term for what happens when the receiving side is not built: people do not enter community life, they move into another institution. The destinations differ by country and by era, and the categories to know are prisons, nursing and old-age homes, custodial or beggar homes, and unregulated private facilities, with homelessness as the outcome when no institution absorbs them at all.

The concept is examinable because it explains why bed-count data can look like success while the population's situation has not improved. A country that halves its psychiatric beds and doubles the psychiatric population of its prisons has not deinstitutionalised anybody; it has changed which ministry is responsible. The corollary for evaluation is that the outcome measure cannot be a psychiatric bed count, and has to be a description of where people actually are.

TRAP **Treating bed reduction as the outcome of deinstitutionalisation.** Bed numbers measure the supply side and are easy to obtain, which is why they get reported. The outcome of interest is where the people are living, whether they are still in contact with services, and whether they have an income and a role. If those cannot be reported, the honest statement is that the programme's inputs are known and its outcomes are not.

23.5 What the Indian record actually documents

India's deinstitutionalisation story is not primarily a story of planned bed closure. It is a story about people who no longer need to be in hospital and are still there, and it has been documented mainly through litigation and inspection rather than through programme evaluation. Three kinds of record exist and each is a different sort of evidence.

The first is judicial. Deinstitutionalisation proceedings before the Supreme Court identified, as a population requiring rehabilitation-home arrangements, persons living with mental illness who have been cured and who do not need further hospitalisation. Note the construction of that category. It is defined administratively, by two negatives, no longer ill in the relevant sense and no longer needing a hospital, and it says nothing about what the person can do or wants. A category defined that way identifies who should not be where they are without identifying where they should go, which is precisely the gap the proceedings were addressing.

The second is inspectorial. **An account by the National Human Rights Commission of visits to 47 mental-health institutions reported about 2,000 cured patients still languishing in institutions**. Read the scope carefully before using the figure. It describes patients in the institutions that were visited, so it is not a national census, it carries no denominator, and no proportion may be computed from it. What it establishes is that the phenomenon exists at institutional scale and has been found by an official body that went and looked.

The third is audit. A performance audit found that a state mental health authority had not conducted the audit of registered mental-health establishments that it was required to conduct. That finding is about the machinery rather than about patients, and it is the most quietly damaging of the three, for the reason set out below.

47	about 2,000	None	Not established
MENTAL-HEALTH INSTITUTIONS VISITED IN THE INSPECTION ACCOUNT	CURED PATIENTS STILL IN THOSE INSTITUTIONS, NOT A NATIONAL CENSUS	VALID DENOMINATOR FOR THAT COUNT, SO NO PROPORTION	WHAT HAPPENS TO PEOPLE AFTER DISCHARGE IN INDIA

23.6 The accountability machinery, and why the audit finding matters

A system that intends to move people out of institutions needs to know what its institutions are doing, which is why regulatory inspection of registered establishments exists as a duty rather than as a good idea. When an audit finds that the required inspection of registered mental-health establishments has not been carried out, the consequence is not merely a procedural lapse. It means that the state's own information about conditions, about length of stay and about who is ready for discharge is not being generated, and a transition cannot be planned on information that is not collected.

This is the general form of a lesson that recurs across this chapter: the failure is usually not at the level of policy, which is generally sound, nor at the level of clinical practice, which is generally willing, but at the level of the machinery that is supposed to observe and report. A resident who can identify the missing observational step in any programme is answering these questions at the right level.

IN INDIA

The operational decision is what to do when you have a patient who is ready for discharge and has nowhere to go. Three things follow from the record above and each is actionable. Record readiness for discharge explicitly in the notes on the day it is reached, because the category of person who no longer needs hospital care is the category the courts and the inspectorate act on, and an undocumented readiness is invisible to both. Involve the social work function early, because the obstacle is an address, an income and a receiving family rather than a treatment. And ask whether your establishment has been inspected by the state authority, because where that machinery is not operating, nothing about your ward's long-stay population is being reported to anybody.

23.7 What India does not know about its own discharges

Here the chapter has to be plain. No official Indian longitudinal or comparative data on what happens to people after discharge from mental-health institutions was retrieved. That includes readmission, sustained community residence, homelessness, incarceration, transfer into custodial or beggar homes, and every other transinstitutionalisation outcome. It was sought from the courts, the health ministry, the human rights commission, the national auditor, the disability empowerment department and state governments.

That absence is not a small one, and it is worth stating what it means rather than merely noting it. Without post-discharge outcome data, nobody can say whether Indian deinstitutionalisation, to the extent that it has happened, has produced community living or transinstitutionalisation. The country has documentation of the problem, judicial direction about it, and no measurement of the result. In a viva, a candidate who says exactly that, and who names transinstitutionalisation as the specific thing that would be measured, has given the strongest available answer.

PITFALL

Assuming that a discharge to family is the same as community living. In a family-based care system the family is the default destination, and for many people it is a good one. But a discharge that transfers the whole burden of supervision, income, occupation and crisis management to a household with no support, no training and no line to the service is a private institution with fewer safeguards than the public one. The question to ask before such a discharge is what the family is being asked to do, and what the service is offering them in return.

23.8 The lessons, stated as design rules

LESSON	WHAT GOES WRONG WHEN IT IS IGNORED	WHAT IT REQUIRES IN PRACTICE
Build the receiving capacity before reducing the sending capacity	People leave with nowhere to go, or do not leave at all and occupy beds that are no longer clinically needed	Housing, income and follow-up commissioned as part of the same programme, not as a later phase
Measure where people are, not how many beds there are	Bed reduction is reported as success while the population moves into other institutions	Post-discharge follow-up with named outcomes, including incarceration and homelessness
Treat readiness for discharge as a documented status	People ready to leave are invisible to the courts, the inspectorate and the planners	Explicit recording, and a register that somebody outside the institution can read
Keep the inspection machinery working	The state stops knowing what its own establishments contain	The statutory authority actually conducting the audits it is required to conduct
Fund the family as a service partner	Discharge becomes an unsupported transfer of custodial responsibility to a household	Carer support, crisis access and a named contact for the family

MNEMONIC

Roof, Rupees, Role, Review

The four things a person needs on the day they leave a long stay. Somewhere to live, an income or entitlement, something to do, and a clinical contact who will notice if they disappear. A discharge plan missing any one of the four is the mechanism by which deinstitutionalisation becomes transinstitutionalisation.

Deinstitutionalisation succeeds or fails on the receiving side, its true outcome measure is where people are living rather than how many beds were closed, and in India the long-stay population has been documented by courts and inspectors while what happens after discharge has never been measured at all.

24 · Community mental health services, day care and halfway homes

Between the ward and the family home there is a ladder, and most Indian districts own only its two ends. A person recovering from a severe illness rarely moves in one step from inpatient care to independent living, and the settings that occupy the middle of that ladder, day care, the halfway home, supported accommodation, longer-term residential care and the community team that visits, exist to make the intermediate steps possible. This section is about how each of those settings is configured: what it is for, who enters it, who staffs it, what a person does there, how they leave, and which arm of government pays for it.

The settings themselves will already be familiar. Their compact definitions, and the place each occupies on the rehabilitation ladder, are taught in the Schizophrenia: management,

antipsychotics and adverse effects notes, and the statutory right to community living and the prohibition on keeping a person in an institution who does not need to be there belong to the Mental health legislation and forensic psychiatry notes. Neither is repeated here. What follows takes the settings as known and asks the service questions about them, which is the part no other chapter carries.

24.1 What it means to configure a setting rather than to define it

A setting is defined by what it is called and configured by five decisions, and it is the five decisions that determine whether it works.

- **Function.** What is the setting actually for, in one sentence: containment, structure, skill acquisition, supervision of treatment, respite for a family, or transition to somewhere else.
- **Entry criteria.** Who is eligible, who decides, and what a person must be able to do before they are accepted.
- **Staffing.** Which disciplines, at what density, present at which hours.
- **Duration and exit.** How long a person stays, what triggers the move on, and where they move to.
- **Funding and governance.** Which arm of government or which body pays, registers and inspects it.

The fourth is the one that separates the settings from each other conceptually, and the fifth is the one that separates them administratively in India, which is why the last two subtopics of this section are about money and counting rather than about clinical content.

RULE **A community setting without a defined exit becomes a small institution.** Every intermediate setting is justified by the step it enables. Where no exit criterion is written and no destination exists, length of stay drifts upward, the population ages in place, and the setting acquires exactly the long-stay character the ladder was built to avoid. Write the exit before the admission.

24.2 The community team, which is the setting without a building

Before any of the physical settings, the ladder needs the service that connects them, and it is the one most often left out of an answer because it has no address. A community mental health team is defined not by where it sits but by where it goes: it holds a caseload of people living at home or in supported settings, sees them where they are rather than where the hospital is, and takes responsibility for noticing when somebody stops attending.

Its four functions are worth naming separately. It supplies *continuity*, so that a person is followed by a service rather than by a series of unconnected appointments. It supplies *outreach*, meaning contact that is initiated by the service, which is the single feature that distinguishes community care from an outpatient clinic and the one that reaches people who default. It supplies *coordination*, holding the relationships with the day-care centre, the residential setting, the family, the employer and the welfare office. And it supplies *crisis access*, without which every deterioration becomes an emergency admission.

In India this function is carried by the district mental health programme's team, and it is worth being clear about the implication. Where that team's posts are vacant or its outreach component is not funded, the entire ladder above and below it becomes a set of disconnected facilities, because there is nobody whose job it is to move a person along it. The physical settings are visible and get built; the connecting function is invisible and gets cut, and its absence is the commonest reason a district with facilities still has no community service.

24.3 Day care, configured

The **day care** centre is the least expensive intermediate setting because it borrows the person's own accommodation. Its function is structure and occupation during the day, together with a regular point of clinical contact, for a person who lives at home. Its natural population is a person whose family can hold the nights and the weekends but cannot supply a day that has any shape to it, which describes a very large share of the severe-illness population in a family-based system.

Its configuration questions are specific. Attendance is usually flexible, so a register that records who is expected and who came is the difference between a service and a room. Transport is the commonest reason for non-attendance and is almost never budgeted. Staffing needs to include somebody who can run an activity programme, not only somebody who can supervise medication. And because the person goes home each evening, the family is a partner in the day care arrangement rather than a bystander, which means the setting can be used deliberately to give a carer predictable hours of relief.

24.4 The halfway home, configured

The **halfway home** is residential and transitional, and both of those words are load-bearing. Residential, because it supplies accommodation and therefore reaches the person for whom home is not currently available or not currently workable. Transitional, because its whole justification is that the person is going somewhere afterwards.

Indian scheme design states the target group explicitly. **The Deendayal Disabled Rehabilitation Scheme of 2009 targeted halfway-home rehabilitation at people whose mental illness is treated and controlled, after their discharge**. Read the eligibility carefully, because it does two things at once. It defines the setting as post-discharge and therefore downstream of treatment, and it defines the population by clinical stability, which means the halfway home is not a place to send an unstable person and not a substitute for admission. A service that uses it as either has misconfigured it.

The configuration questions that decide whether such a home works are the length of stay and the destination. A halfway home with no ceiling on stay and no supported accommodation to move into becomes long-term residential care under another name, which is not a failure of the staff but a failure of the ladder above them.

TRAP **Describing the halfway home as a place for people who are unwell.** The scheme design defines its population as people whose illness is treated and controlled, after discharge. Candidates routinely describe it as an intermediate treatment setting for partially recovered patients, which imports an acute function it was not configured for and confuses it with day care and with continuing residential care in the same answer.

24.5 The continuing-care tier, and how India planned it

Not everybody moves on. A proportion of people with the most severe and persistent illness, particularly those with no family and no income, need residential support that is not transitional, and a ladder that pretends otherwise simply discharges them into homelessness. Indian programme design has recognised this in the structure of its own planning: **the National Health Mission implementation framework for 2012 to 2017 provided for day-care centres, residential continuing-care centres and a long-term residential continuing-care centre in selected districts**.

Three features of that plan are worth teaching. It is tiered, distinguishing day attendance from residential continuing care and again from long-term residential continuing care, which is exactly the ladder logic. It is explicitly limited to *selected* districts, so it was never a universal offer and should not be described as one. And it is a statement of what will be provided, which is a plan rather than an inventory, so nothing about how many such centres came to exist follows from it.

24.6 Two doors, two ministries, one patient

The single most useful administrative fact in this topic is that India's intermediate settings are reachable through two different arms of government, and they have different eligibility logics, different registering authorities and different money. The health route runs through the national and district mental health programmes and treats the setting as a health service. The disability and social welfare route runs through disability schemes and treats the same setting as a rehabilitation facility for a person with a disability, which brings certification, entitlements and a different inspection regime with it.

ROUTE	HOW THE SETTING IS FRAMED	TYPICAL ELIGIBILITY LOGIC	WHAT IT BRINGS WITH IT
Health programme route	A service component of the mental health programme, planned for selected districts as day care and as two tiers of residential continuing care	Clinical need, decided by the treating service	Clinical staffing, a link to the district team, and dependence on programme funding cycles
Disability and social welfare route	A rehabilitation project, with the halfway home as a named model project for people whose illness is treated and controlled after discharge	Post-discharge stability, and often a disability certificate	Grant-based funding to voluntary organisations, certification requirements, and a different inspectorate

The practical consequence for a resident is that the answer to "where can I send this person" depends on which door you knock on, and that the two doors have different waiting times, different paperwork and different reasons for refusing. Knowing that there are two is worth more in a district job than knowing the name of any single facility.

IN INDIA

The decision this section forces is procedural and it is made early or not at all. If a person is likely to need an intermediate residential setting after discharge, start the disability certification and the scheme paperwork while they are still an inpatient, because the social welfare route requires documentation that takes time to assemble and cannot be produced on the day of discharge. Simultaneously ask the district programme what it actually operates locally, since the health route's day-care and continuing-care provision was planned for selected districts rather than for all of them. Doing both in parallel is the only way to have a destination ready when the person is ready.

24.7 What is counted, and what a count of it means

A national dashboard publishes a count of rehabilitation homes and halfway homes across states and union territories, which sounds like the coverage figure this section needs and is not. This chapter cannot print those figures, because no claim row here asserts them, and the honest treatment is to teach what such a count would and would not establish.

A dashboard count is a count of *listed or registered* facilities. It is not a count of facilities that are open, not a count of places available, not a count of people accommodated, and it has no denominator, so no proportion of need met can be computed from it. Two facilities appearing as one row each may differ by an order of magnitude in the number of people they hold. And a national total says nothing about distribution, which is the variable that matters to a patient, because a home in another state is not a destination.

The deeper gap sits underneath the count. Current operative staffing standards and referral criteria for day-care centres and halfway homes, together with any national or state outcome data that would distinguish sanctioned or registered provision from functioning services, were sought across the health ministry and mission, the disability empowerment department and state health and social welfare authorities, and were not retrieved. So the position to state in a viva is this: the settings are named in Indian scheme design, their target populations are defined, a national listing of homes exists, and what they are staffed with, who they must accept and what happens to the people in them is not established.

PITFALL

Quoting a registry total as service capacity. It is the same error as reading a sanctioned district total as coverage, and it recurs across this chapter because registries are the only national data that exist for community services. A listed home may be closed, may be full, may take only one sex, may charge fees, and may be a day's travel away. Before referring, confirm the four things a registry never records: is it open, is there a place, what does it cost, and will it take this person.

Supported accommodation deserves a mention as the rung above the halfway home, because it is the destination that makes a transitional setting transitional. It is ordinary housing with a support component attached, ranging from a warden or visiting worker to a shared flat with staff on call, and its defining feature is that the tenancy rather than the treatment is the organising fact. Where it does not exist, everything below it silts up. This chapter's sources carry no figure for its availability in India, so the honest statement is that the rung is named in policy and its extent is unmeasured.

24.8 Choosing a setting for a particular person

The decision reduces to three questions asked in order. Does the person have somewhere to sleep tonight that is safe and available? If yes, the residential settings are not indicated and the question is what fills the day. If no, is the need transitional or continuing? A transitional need with a plausible destination points to a halfway home; a continuing need with no destination points to the continuing-care tier, and pretending otherwise produces a failed placement. And finally, who is going to keep clinical contact wherever the person goes? A community setting without a named clinical contact is where people are lost to follow-up.

MNEMONIC

Function, Entry, Staffing, Exit, Payer

The five configuration questions, in the order in which they should be asked of any community setting. If you can answer all five for the day-care centre and the halfway home in your own district, you know more about local services than most of the people referring into them.

Selected DISTRICTS THE PLANNED DAY-CARE AND CONTINUING-CARE TIERS WERE FOR	**Post-discharge** TARGET POPULATION FOR THE HALFWAY-HOME MODEL PROJECT	**Two** GOVERNMENT ROUTES INTO THE SAME SETTINGS, HEALTH AND SOCIAL WELFARE	**Not established** STAFFING STANDARDS, REFERRAL CRITERIA AND OUTCOMES FOR EITHER SETTING

Day care borrows the person's own accommodation and supplies the day; the halfway home is residential, transitional and configured for people whose illness is treated and controlled after discharge; and in India both are reachable through a health route and a social welfare route with different eligibility, different money and different inspectors.

25 · Models of care in India: NGOs, family-based care and indigenous resources

Most mental health care in India is delivered by people the health system does not employ. Families, voluntary organisations, community volunteers and traditional practitioners between them carry a far larger share of the load than the district programme does, and a chapter on service architecture that describes only the government machinery has described the smaller part of the system. This section treats those three as *models of care*: configured arrangements with an identifiable provider, a funding basis, a governance structure and, in principle, an evaluable output.

Two boundaries apply before we start. The family as the delivery vehicle for psychoeducation and for work with high expressed emotion, and the voluntary sector named as a force multiplier for community reach in the management of schizophrenia, are both taught in the Schizophrenia: management, antipsychotics and adverse effects notes, and neither is restated here. The design of a formal partnership with traditional and faith-based healers, meaning bidirectional referral, training, safeguards and access, is taught in the section of this chapter that owns traditional healers, and is named here rather than built.

25.1 What makes something a model rather than an activity

A **model of care** is not the same as a thing that helps. Four properties turn an arrangement into a model, and they are the properties that let it be commissioned, replicated and evaluated.

- **A defined provider and a defined role.** Somebody specific does something specific for somebody specific, and it is written down.
- **A funding and governance basis.** Money comes from somewhere on a stated basis, and somebody is accountable for how it is used.
- **A relationship with the formal system.** Referral in, referral out, and an agreed division of what each side does, without which the model is a parallel service rather than a component of one.
- **An output that could be measured.** Not necessarily measured, but of a kind that could be.

Running the three Indian limbs through those four properties is the most useful thing this section does, because it explains immediately why the voluntary sector is the best documented, why family care is the largest and least visible, and why indigenous resources are the hardest of the three to hold to any of it.

RULE **Name the provider, the money and the referral relationship, or you have not described a model.** A great deal of writing about Indian community psychiatry describes admirable activity without any of the three, which is why so little of it can be commissioned by anybody or repeated anywhere else.

25.2 The voluntary sector as a configured model

The **voluntary sector**, meaning non-governmental organisations working outside both the state system and private practice, occupies a specific and predictable position in Indian mental health care. They tend to do the things the public system finds hardest to fund: long-term residential rehabilitation, day activity, vocational work, homeless outreach, carer support, and services for populations that fall between departmental mandates. They are usually geographically concentrated, so their national coverage is thin even where their local quality is high, and they typically depend on a mixture of philanthropic funding, government grants and project money, which means their planning horizon is set by their funding cycle rather than by their patients' needs.

The partnership questions between such an organisation and a district service are concrete and are what a resident actually has to negotiate. Who accepts clinical responsibility for a person in

the organisation's care, and is that written down? What is the referral route in each direction, and what happens outside working hours? Who holds the clinical record, who prescribes, and who reviews? What happens to the people in the service if the grant ends? An arrangement with clear answers is a model; an arrangement without them is a friendship between two institutions, and it does not survive a change of personnel.

This section names no Indian organisation. That is a deliberate omission and worth stating plainly: naming a service is a factual claim about what that organisation does, and this chapter's sources do not carry a description of any individual Indian organisation. A resident preparing for an examination should know the sector's characteristic functions, which are set out above, and should read about specific organisations in their own published accounts rather than from a list learned second-hand.

25.3 The family as India's largest care system

In India the family is not an adjunct to care but the substrate on which everything else rests. Co-residence is common, supervision of medication is routine, transport to appointments is a family undertaking, and the family is usually the first to detect deterioration and the last line of support when nothing else is available. Reading that as a service model rather than as a cultural fact changes what a service is obliged to do about it.

If the family is a provider, then it has the same needs as any other provider. It needs to know what it is treating, which is information rather than instruction. It needs a route of access when things change, because a carer with no way to reach the service will present at an emergency department instead. It needs relief, since supervision without respite is not sustainable across the years these illnesses last. It needs financial support, because carer time is unpaid labour with an opportunity cost that falls disproportionately on women. And it needs somebody to notice its own health, since carer distress is both a consequence of the role and a determinant of whether the person being cared for stays at home.

Carer organisations and family self-help groups are the organised form of this. Their functions are advocacy, mutual support, and in some settings direct service provision, and they change the relationship between families and services from petitioning to negotiating. The techniques of family intervention itself, psychoeducation and the work on expressed emotion, belong to the notes cross-referred above and are not described here; what belongs here is the service obligation that follows from treating the family as part of the workforce.

TRAP **Describing family involvement as an Indian cultural advantage and stopping there.** It is an advantage, and it is also an unfunded transfer of long-term care to households, disproportionately to women, with no training, no relief and no income attached. Both halves belong in the answer. A candidate who writes only the first half has described a resource; a candidate who writes both has described a model with a cost.

25.4 Indigenous and traditional resources as a service model

Indigenous resources, the third limb, are what no other chapter treats as a service model, and they are difficult precisely because the phrase covers at least three different things. There are

traditional systems of medicine with their own texts, training and practitioners. There are religious and faith-based settings to which distressed people and their families go, sometimes in very large numbers. And there are local healers with no formal system at all, who are frequently the first point of contact and, in many districts, the most accessible one.

Treated as a service model, the honest ledger is short and should be stated without either romanticism or contempt. What these resources reliably supply is *access*, because they are local, affordable, culturally intelligible and available without a referral; *explanation*, in an idiom the family already uses; and *social sanction*, since consulting them carries none of the discredit a psychiatric attendance carries. What they do not supply is evidence of therapeutic effect for psychiatric disorder, and this chapter does not print one. The cultural and relational significance of indigenous healing, including the traditional teacher and disciple relationship read as a native model of the therapeutic bond, is taught in the Psychotherapies, clinical assessment, MSE and nosology notes, which places it at the cultural level and explicitly not as evidence-based treatment. This chapter prints no claim stronger than that.

The service consequence is a matter of position rather than of belief. These practitioners are already in the pathway, whether or not any service acknowledges them, so the operative question is not whether they should be, but what a formal service does about the fact. The design of that relationship, referral in both directions, what may safely be shared, what training is offered and what safeguards apply, is the subject of the traditional healers section of this chapter and is not settled here.

25.5 The one Indian community model with a trial behind it

Very few Indian community models have been evaluated with a controlled design, which is precisely why the ones that have been matter disproportionately. A community-champion model deserves attention on those grounds: in a cluster-randomised trial conducted in rural Gujarat, **the Atmiyata intervention had a significant effect on recovery from symptoms of depression and anxiety, with effects sustained at 8-month follow-up**.

What makes this instructive is the provider. The intervention is delivered by volunteer community champions, ordinary community members rather than health workers, which places it at the far end of the workforce spectrum from the specialist service. Three features of the design are worth extracting for any examination answer. The provider is drawn from the community being served, which addresses both access and acceptability at once. The design is cluster-randomised at village level, which is the right unit when the intervention acts on a community rather than on an individual. And the follow-up extends beyond the intervention period, which is what turns a demonstration of effect into a claim about durability.

IN INDIA

The decision this section supports is how to treat the non-governmental providers already working in your district. Treat them as a model rather than as goodwill: ask each one the four questions that make a model, who does what, on what money, with which referral route in each direction, and what they count. Where a volunteer or community-champion layer exists, the Indian evidence supports taking it seriously as a therapeutic component with a durable effect rather than as outreach, which means it deserves supervision, a defined scope and a route to escalate, exactly as a paid worker would.

25.6 What the Indian model literature is missing

The gap in this topic is not a gap in activity but a gap in evidence and in accounting, and it has a characteristic shape. There is no national inventory of voluntary-sector mental health provision, so nobody knows what exists or where. There is no measurement of what family care actually costs the households that supply it. Very few community models have been evaluated against any comparator, and those that have are concentrated in a handful of research settings, which raises the transportability question rather than settling it. And the most common form of publication is a description of a promising programme with no control condition and no follow-up, which is why so much of the field reads as encouraging and proves so little.

Because so much of the literature is descriptive, it is worth carrying a short appraisal routine for any Indian model-of-care paper, and it is the routine an examiner is testing when they ask what you think of a programme. Ask what the comparator was, and treat "before and after in the same people" as the weakest answer available, because a service that recruits people at their worst will improve without doing anything. Ask who was counted, and specifically whether the denominator is everybody offered the model or only those who stayed in it. Ask when the outcome was measured, since an assessment taken at the end of the intervention says nothing about durability. Ask who delivered it and whether that person exists outside the study, because a model delivered by a research team is not a model a district can run. And ask what it cost, since a model whose cost is unreported cannot be compared with anything the state is already funding.

Applying that routine to the Indian field produces a consistent verdict: a small number of models are properly evaluated, a large number are described, and the two groups are not the same models. That is not a criticism of the people running services, who are rarely funded to evaluate anything. It is an argument for building the evaluation question into the design of the next programme rather than attempting to reconstruct it afterwards.

None of that is a reason to dismiss the models. It is a reason to be precise about what is claimed for each, which is what the table below does.

MODEL	WHAT IT RELIABLY SUPPLIES	FUNDING AND GOVERNANCE	WHAT IS ESTABLISHED ABOUT EFFECT
Voluntary sector organisations	Long-term rehabilitation, day activity, vocational work, outreach to populations the public system misses	Philanthropy, grants and project money, with a planning horizon set by the funding cycle	Nothing at national level in this chapter's sources; no national inventory of provision exists
Family-based care	Supervision, accommodation, transport, early detection and the whole of long-term support	Unfunded, borne by households and disproportionately by women	Nothing quantified here; the cost to households is unmeasured
Community champions and volunteers	Local, acceptable, first-line psychosocial contact within the community itself	Programme or research funding, with volunteers rather than salaried staff	A cluster-randomised trial in rural Gujarat found improved recovery from depression and anxiety symptoms, sustained on follow-up
Indigenous and traditional resources	Access, explanation in the family's own idiom, and social sanction	Private, fee-for-service or offering-based, outside any health governance	No effect claim is printed here; the cultural reading is owned by another chapter and is explicitly not an evidence claim

PITFALL

Treating an unevaluated local programme as a model because it is admired. Enthusiasm, longevity and local reputation are not evidence, and a programme that has run for years without ever defining an output is not a model that anybody else can adopt. The kindest and most useful thing a visiting psychiatrist can do for such a programme is to help it define what it would count, because that is what converts it from a good thing into something a state could fund.

MNEMONIC

Provider, Purse, Pathway, Proof

The four questions that turn an arrangement into a model of care. Who does it, who pays for it, how does a person get in and out of it, and what would show that it worked. Ask them of a voluntary organisation, of a family and of a village volunteer alike, and the differences between the three limbs become clear immediately.

Volunteers
PROVIDER IN THE ONE INDIAN COMMUNITY MODEL WITH A CONTROLLED TRIAL HERE

Sustained
RECOVERY EFFECT AT FOLLOW-UP IN THAT TRIAL

Unfunded
BASIS ON WHICH FAMILY CARE IS SUPPLIED

No inventory
NATIONAL RECORD OF VOLUNTARY-SECTOR MENTAL HEALTH PROVISION

A model of care has a provider, a purse, a pathway and a proof; India's three great non-governmental models are the voluntary sector, the family and the indigenous practitioner, and of the three only a community-volunteer intervention carries a controlled Indian trial in this chapter's evidence.

Primary care, integration and service delivery models

26 · Integration of mental health into primary care and the role of medical officers

There will never be enough psychiatrists, and integration is the only answer anybody has that is not arithmetic nonsense. If mental disorder is common and specialists are few, then care has to be delivered by the part of the health system that is already everywhere and already sees everybody, which is primary care. Integration is the deliberate rebuilding of primary care so that it detects, treats, refers and follows up mental, neurological and substance-use conditions as ordinary business rather than as a specialist favour. This section is about how that is designed in India, who does what, and what is known about whether it is happening.

Two internal boundaries keep this section honest. The competencies, training and supervision that make non-specialist delivery possible, together with the decision-support instrument built for it, belong to the section of this chapter on task sharing. The escalation architecture of collaborative and stepped care, meaning who is added to a person's care and when, belongs to the section that owns those models. The national programme's own objectives and evolution belong to the section that owns the national programme, and are named here rather than restated.

26.1 What integration is, and three things it is not

Integration means that a mental health task is performed at the primary level by the staff who are there, using the same record, the same clinic session and the same supply chain as everything else they do. Being clear about what it is not is the faster route to the concept.

- **It is not a specialist clinic held in a primary-care building.** A visiting psychiatrist running a monthly session at a health centre is outreach. It is useful, it increases access, and it leaves the primary-care team's own practice exactly as it was, because nothing is done differently on the other twenty-something working days of the month.
- **It is not a screening programme.** Detection without a treatment response and a follow-up route generates identified and untreated people, which is worse than not looking, because it consumes trust as well as time.
- **It is not the transfer of all mental health care to primary care.** Integration implies a division of labour, and a division of labour implies that some conditions and some phases go upward. A model with no referral criteria has not integrated anything, it has abandoned people at the bottom of the system.

RULE **Integration is a two-way relationship or it is not integration.** Every design in this field has an upward limb, referral of what primary care cannot manage, and a downward limb, return of the person to primary care with a plan the primary-care team can execute. Systems build the upward limb because it is a request and neglect the downward limb because it is a gift. The result is a specialist service that fills with people who could be managed below it.

26.2 The medical officer's workflow, which is the heart of the item

Indian operational guidance is explicit about what the doctor at the primary level is expected to do. **At the primary and urban primary health centre level the medical officer's clinical workflow spans case detection, management, referral and follow-up of mental, neurological and substance-use disorders**. Those four verbs are the item in miniature, and each of them has a distinct requirement attached.

Detection requires that the presenting complaint be recognised for what it is, which in primary care is very often a bodily symptom, a sleep complaint, a request for a tonic, or a family member's description of behaviour rather than a patient's description of mood. It also requires time, which is the scarcest commodity in a busy general clinic.

Management requires a small formulary that is actually stocked, a small set of conditions the officer is expected to treat, and permission to treat them. Where any of the three is missing the workflow stops at detection.

Referral requires a named destination, a route and a reason to believe the person will be seen. A referral into a queue nobody can enter is a form of discharge.

Follow-up requires a record that persists between visits and somebody who notices non-attendance, which together are what the field calls **continuity of care**. This is the step most often absent, and its absence converts a chronic-disease service into a series of unrelated first visits.

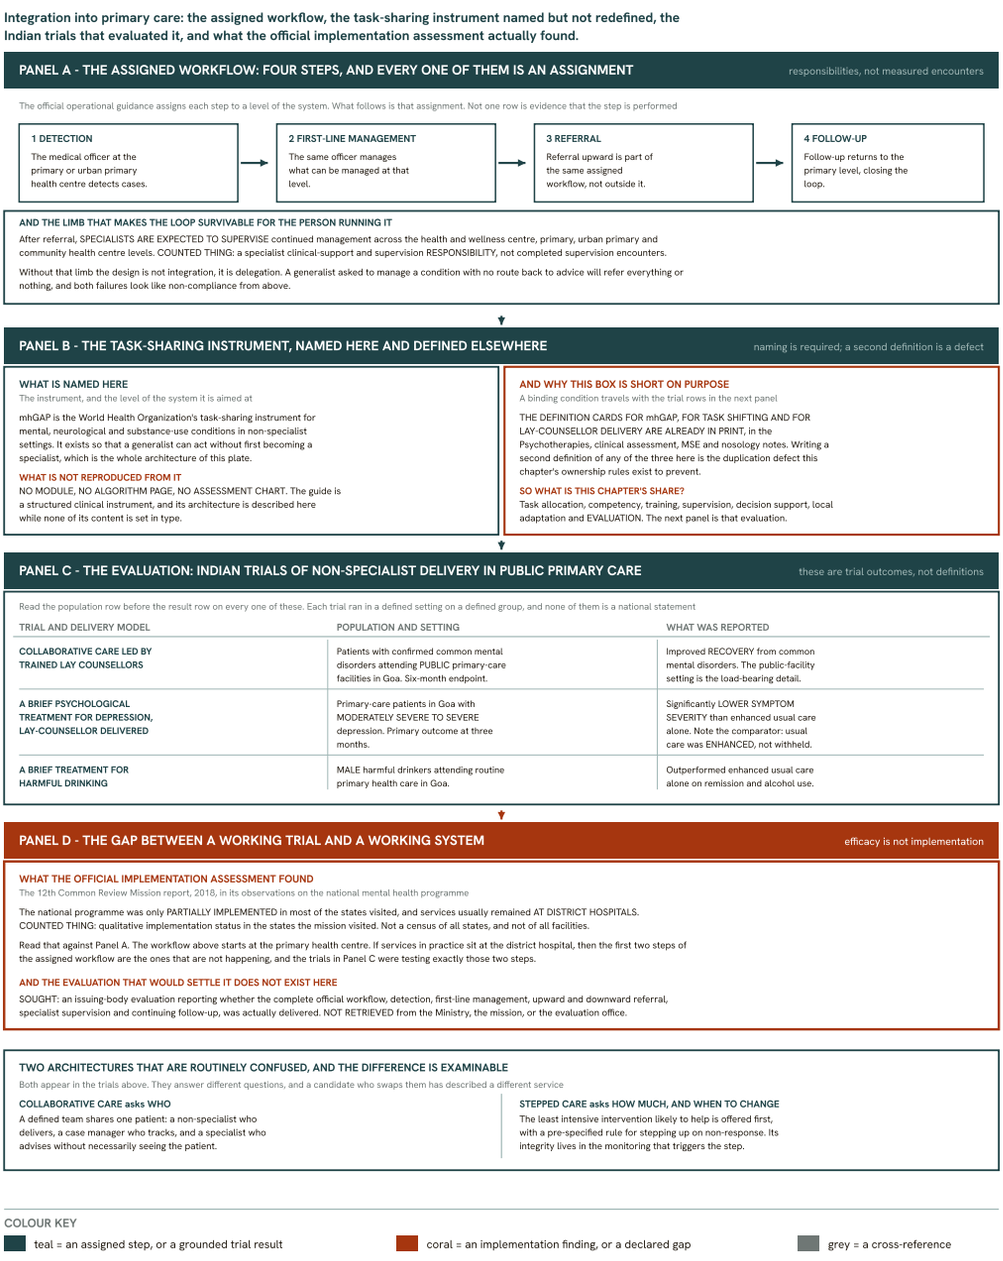


Fig 39.9 Primary-care integration: the assigned workflow, the instrument named, the trials that evaluated it, and the implementation finding that sits against all three. The medical officer at the primary or urban primary health centre is assigned case detection, management, referral and follow-up, and specialists are expected to supply continuing clinical support and supervision after referral, which is what separates integration from delegation. mhGAP is named here as the task-sharing instrument; its definition, and those of task shifting and lay-counsellor delivery, are already in print in the Psychotherapies, clinical assessment, MSE and nosology notes and are not rewritten. The evaluation is Indian and it is real: lay counsellor-led collaborative care improved recovery in public primary-care facilities, a brief lay-delivered treatment lowered depression symptom severity against enhanced usual care, and a counterpart for harmful drinking outperformed enhanced usual care alone. Against that, the official implementation assessment found the programme only partially implemented, with services usually remaining at district hospitals, and no issuing-body evaluation of the complete workflow was retrieved.

26.3 The upward limb, and the return leg that makes it a loop

The Indian programme design for severe illness is specific and worth learning as a sequence. **Within the national and district programme's service delivery, primary health centres refer people with severe mental disorders to the district hospital, and then follow up the treatment plan drawn by the psychiatrist there**. Read the second half of that sentence carefully, because it is the part that gets lost. The primary health centre is not simply a referring agency; it is the place where the specialist's plan is subsequently executed, month after month, which is where the great majority of the person's contact with the health system actually occurs.

Three practical requirements follow from the return leg and each is a common point of failure. The plan has to be written in a form the primary team can act on, which means named drug, dose, duration, what to watch for and when to send the person back. The medicine has to be available at the primary level, since a plan that requires a monthly journey to the district hospital for a supply has not returned the person anywhere. And the route back up has to be open without a fresh referral for a person who deteriorates, because a rigid one-way system punishes the primary team for accepting the patient in the first place.

TRAP **Describing integration as detection plus referral.** This is the commonest incomplete answer in the topic. Detection and referral are the first and third of four verbs in the official workflow, and a description that omits management at the primary level and follow-up after the specialist visit has described a triage service. The examinable content is the loop, not the arrow.

26.4 Supervision, which is what holds the whole design up

Supervision, meaning a continuing relationship in which a non-specialist can bring a real case to a specialist, is the component everything else rests on. A non-specialist asked to manage a condition they were not trained to manage will do it once, badly, and then stop, unless somebody is available to be asked. Indian guidance places that obligation explicitly: **after referral, specialists are expected to provide ongoing clinical support and supervision for the continued management of people with these conditions at the community health centre, primary health centre and health and wellness centre levels**.

Note what this is not. It is not a training event, which is a discrete input; it is a continuing relationship. Note also who it burdens: the same district psychiatrist who is already running the outpatient department is the person expected to supply it, which is why supervision is the component that quietly disappears when a post falls vacant. The design questions that decide whether it survives are mundane and answerable: what is the medium, is it scheduled or on demand, is it case-based or didactic, and is anybody's job description changed to include it. The competency framework and training architecture that sit behind supervision belong to the task-sharing section of this chapter.

IN INDIA

The decision this section supports is what to do with the referral letter you write as a district psychiatrist. Because the primary health centre is expected to follow the plan you draw, the letter is not a discharge summary but an operating instruction for a colleague who has no specialist training and no easy way to ask you a question. Write the drug, the dose, the intended duration, the two adverse effects that would matter, the specific circumstances that should bring the person back to you, and a way to reach you. A plan that assumes the primary team will interpret it is a plan that will be abandoned at the first uncertainty, and the person will simply stop attending.

26.5 What is actually happening, as opposed to what is designed

India has operational guidance, a defined workflow, a referral pathway and a supervision obligation. What it has less of is evidence that any of this is running. The most useful official observation available here is a review-mission finding: **the 2018 Common Review Mission found the national mental health programme implemented only partially by most of the states visited, with services usually available through district hospitals**.

That single sentence contains the central problem of this item. If services are usually available *through district hospitals*, then care has not descended to the primary level in most of the places that were looked at, which is the opposite of integration. Two cautions on scope: the finding describes states and facilities that a review mission visited rather than a census of the country, and it is qualitative, so it supports a direction and not a proportion.

The evaluation that would settle the question does not exist in this chapter's sources. No issuing-body evaluation reporting whether the complete official workflow, detection, first-line management by the medical officer, referral upward and downward, specialist supervision and continuing follow-up, has been implemented in Indian primary-care facilities was retrieved. It was sought from the health ministry, the national health mission, the national health systems resource centre and the government's evaluation office, and what was found was either an earlier review, a document predating the current operational guidance, or a procurement notice for an evaluation rather than its findings.

Four	District hospital	Partial	Not retrieved
VERBS IN THE OFFICIAL PRIMARY-CARE WORKFLOW: DETECT, MANAGE, REFER, FOLLOW UP	WHERE SERVICES WERE USUALLY AVAILABLE IN THE STATES A REVIEW MISSION VISITED	IMPLEMENTATION FOUND IN MOST STATES VISITED, QUALITATIVE AND NOT A CENSUS	ANY EVALUATION OF THE COMPLETE WORKFLOW IN INDIAN PRIMARY CARE

26.6 Where integration breaks, in the order it breaks

The failure points are predictable, and a resident who can name them in sequence can both diagnose a local service and answer the question.

STEP	WHAT HAS TO BE TRUE	HOW IT FAILS	WHAT WOULD SHOW IT
Detection	Time in the consultation, and recognition of presentations that are not complaints of mood	The presentation is treated as the physical complaint it arrived as, and nothing further happens	Recorded diagnoses at the primary level, against expected frequency
Management	A stocked formulary, a defined scope of practice and permission to prescribe within it	Detection without treatment, so the officer stops detecting	Drug availability at the primary level, and prescriptions actually issued there
Referral upward	A named destination that will see the person	Referral into an inaccessible queue, which the family experiences as refusal	Proportion of referrals that result in a specialist contact
Return downward	An executable plan and the medicines to execute it	The person keeps returning to the district hospital for routine review, which silts up the specialist clinic	Proportion of stable people reviewed at the primary level rather than at the hospital
Supervision	A scheduled, case-based relationship with somebody's job description changed to include it	It becomes voluntary, then occasional, then nothing, especially when a post is vacant	Scheduled supervision contacts held, not training events delivered
Follow-up	A persisting record and somebody who notices non-attendance	Every visit becomes a first visit and continuity is lost	Retention in care over a defined period, with the denominator stated

PITFALL

Judging a primary-care mental health service by how many people it detected. Detection is the cheapest number to produce and the easiest to increase, and on its own it says nothing about whether anybody was treated or stayed in treatment. The number that reflects an integrated service is retention: how many of the people started on treatment at the primary level are still in care after a defined period, with the denominator stated. If a service cannot produce that, it does not yet know whether it works.

Reading the table upward rather than downward gives the other examinable lesson. A service that is failing at follow-up cannot be fixed by more training in detection, and a service failing at management because nothing is stocked cannot be fixed by supervision. Each step depends on the ones before it, so the correct first question about any underperforming integrated service is which step is the earliest one that is broken. Programmes that repeat detection training every year while the formulary remains empty are extremely common, and they are not failing through incompetence; they are fixing the step that is easiest to fix.

26.7 The workforce question, and why this section prints no benchmark

The argument for integration is usually opened with a workforce statistic: so many psychiatrists for so many people, against a recommended minimum. This section prints no such benchmark, and the reason is worth teaching because it is a lesson about evidence rather than about workforce.

Several psychiatrist-per-population figures circulate in Indian teaching, they differ from one another, and the recommended minimum they are usually measured against cannot be traced to an issuing body in this chapter's sources. A ratio quoted without a source, a reference year and a stated definition of who counts as a psychiatrist is not a fact; and a minimum attributed to an international body that did not issue it is a false attribution, which is a more serious error than a wrong number because it borrows authority as well as data. The defensible position is to make the argument structurally, which loses nothing: specialists are few and unevenly distributed, mental disorder is common and present in every district, and no plausible expansion of specialist training changes that arithmetic within a working lifetime. That argument does not need a ratio, and it cannot be wrong.

GOOD TO REMEMBER

When a primary-care colleague telephones about a patient, the useful reply is not a diagnosis but a decision they can act on today: what to start or change, what to watch for, when to review, and what should bring the person back to you. Answering at that level takes less of your time than a referral does, keeps the person near home, and is the single most effective piece of integration available to an individual psychiatrist regardless of what the programme is doing.

Integration is a loop, not an arrow: detect, manage, refer upward and take the person back with a plan the primary team can execute, held together by continuing specialist supervision, and Indian review findings indicate that services in most states visited still sat at the district hospital rather than below it.

27 · WHO mhGAP intervention guide and task-shifting

If the specialist workforce cannot grow fast enough, the only remaining variable is who is allowed to do what. Task sharing is the deliberate redistribution of defined clinical tasks to workers with less training, supported by structured decision aids, supervision and a clear escalation route. The World Health Organization's mhGAP Intervention Guide is the instrument built to make that redistribution safe for mental, neurological and substance-use conditions. This section treats both as design problems: how tasks are allocated, what competency means, how training and supervision are constructed, what a decision aid actually does, how much a country may change, and what the evidence shows.

The definitions themselves are owned elsewhere and are named rather than rebuilt. Task-shifting and task sharing as constructs, the mhGAP Intervention Guide as a document, and lay-counsellor-delivered treatment as a category are all defined in the Psychotherapies, clinical

assessment, MSE and nosology notes, which carries them with their evidence base. The primary-care workflow into which any of this is delivered belongs to the section of this chapter on integration, and the escalation architecture of stepped care to the section that owns it. No module, algorithm or assessment chart from the guide is reproduced anywhere in this chapter.

27.1 Task allocation is a design decision, and it is made task by task

The question is never whether a cadre is "capable". It is whether a specific task, defined narrowly enough to be trained and checked, can be performed by a specified worker under specified conditions with an acceptable failure rate. Four properties make a task a good candidate for redistribution, and the same four explain why some tasks resist it.

- **The task is bounded.** It can be described as a procedure with a start, a finish and a small number of decision points, rather than as clinical judgment applied to an undefined situation.
- **The decision rule is explicit.** Somebody can write down what to do in each case that arises, which is precisely what a decision aid is for.
- **The failure is visible.** If the task is done badly, somebody notices, either because the person deteriorates in a recognisable way or because a check exists.
- **An escalation route exists.** The worker can hand the situation upward quickly when it exceeds the boundary, and doing so is treated as correct practice rather than as failure.

Applying those four in order settles most disputes. Delivering a manualised brief psychological treatment for a defined condition scores well on all four. Adjusting an antipsychotic in a person with a complex physical comorbidity scores badly on the first two. Recognising that somebody is unwell and needs to be seen scores well on all four and is therefore the task most widely and most safely redistributed anywhere in the world.

RULE **Task sharing redistributes a task with its supports, or it is simply understaffing.** The package is training plus a structured decision aid plus supervision plus an escalation route. Remove any one of the four and the model becomes a less qualified person doing a harder job alone, which is not what any of the evidence tested.

27.2 Competency rather than cadre

Competency, meaning what a worker can be observed to do rather than what they hold a certificate in, is the organising idea of this whole section. The professional reflex is to ask what qualification somebody holds. The task-sharing question is what they can demonstrably do, which is a different and more useful thing to know. The practical consequence is that programmes specify competencies, then select whoever can acquire them, which is how community health workers, nurses, counsellors without a psychology degree and general medical officers all end up delivering parts of the same programme in different settings.

Selection matters more than it looks. Programmes that recruit for attributes the training cannot supply, literacy sufficient for the materials, standing in the community, the ability to hold a conversation about distress without flinching, and a willingness to be supervised, achieve competency faster than programmes that recruit on formal qualification alone. And competency

has to be assessed rather than assumed: attendance at a training course is a record of presence, and observed performance on the task is a record of competence.

27.3 Training, and why most of it decays

Training in this field fails in a characteristic way. A cascade model, in which a small group is trained centrally and each trains the next tier, loses content at every step; a single workshop produces a measurable rise in knowledge that decays within months; and knowledge gained in a classroom transfers poorly to a clinic without practice under observation.

What survives has a recognisable shape: it is short, repeated rather than single, built around cases rather than around lectures, includes observed practice, and is followed by continuing contact with the trainer. That last element is the bridge from training into supervision, and it is why the two should be commissioned as one thing rather than as a course and a hope.

27.4 Supervision, the component that is always cut first

Supervision in a task-sharing programme is not the periodic review of a subordinate. It is the mechanism by which a bounded task stays bounded: the point at which cases that have drifted outside the worker's remit are identified and moved, at which errors are corrected before they become habits, and at which the worker is retained, because isolated non-specialists doing emotionally demanding work leave.

Its design parameters are few and are worth being able to state: frequency, who supervises, whether it is individual or group, whether it is case-based or general, and whether attendance is protected time or personal generosity. Group case-based supervision at a defined interval, delivered by somebody one step up rather than by the most senior person available, is the arrangement that most often survives contact with a real service, because it is affordable and because it uses the tier immediately above rather than the scarcest tier.

TRAP **Presenting task sharing as a solution to a shortage of specialists.** It redistributes specialist time rather than removing the need for it. Every task-sharing model consumes specialist time in training, in supervision and in receiving escalated cases, and a district with no specialist at all cannot run one properly. The correct framing is that task sharing multiplies the reach of the specialist who exists, not that it substitutes for one who does not.

27.5 What a decision-support instrument does

The **mhGAP** Intervention Guide is the instrument the field is built around, and its function is worth describing precisely because the description is what an examiner wants. Architecturally it is organised by presentation rather than by diagnosis, it converts assessment and management into stepwise sequences a non-specialist can follow, it specifies what must be excluded before proceeding, and it states explicitly at which points the person should be referred or discussed with a specialist. In other words it externalises the parts of clinical reasoning that ordinarily live in a trained clinician's head, so that the worker's judgment is required in fewer places and those places are marked.

Its limits follow from the same design. A presentation-based sequence handles the common and the typical and is not built for the atypical, the comorbid or the treatment-resistant, which is exactly why the referral points are part of the instrument rather than an afterthought. And a guide is not a programme: possessing it changes nothing without training, supervision and a service to escalate into. This chapter names the guide, describes what it does, and reproduces no part of its content.

27.6 Local adaptation, and what may not be changed

Every country adapts, and adaptation is expected rather than tolerated. The distinction to hold is between adapting the *surface* and adapting the *structure*. **Surface adaptation** covers language, the idiom in which symptoms are described, the examples used, the cadre selected, the setting, and alignment with the national formulary and referral map, and all of it improves fidelity to intent. **Structural adaptation**, meaning removing a referral trigger, widening the task boundary, dropping a step because it is inconvenient, or omitting supervision because there is nobody to supply it, changes the intervention into something the evidence does not cover.

The examinable formulation is that adaptation should preserve the safety architecture and change everything else. A programme that has translated the materials, chosen a locally appropriate cadre and aligned the drugs with what the district actually stocks has adapted well. A programme that has kept the materials in their original form and quietly abandoned supervision has adapted badly while appearing faithful.

27.7 The Indian evidence, which is unusually strong

India contributed two of the most-cited randomised evaluations in this field, both testing brief psychological treatments delivered by lay counsellors in routine primary care against enhanced usual care. For depression, **participants receiving the Healthy Activity Program in addition to enhanced usual care had significantly lower symptom severity than those receiving enhanced usual care alone**, in primary-care patients with moderately severe to severe depression. For harmful drinking, **Counselling for Alcohol Problems delivered by lay counsellors in addition to enhanced usual care was better than enhanced usual care alone for harmful drinkers in routine primary health-care settings**.

Four features make this pair worth learning rather than merely citing. The comparator in both is enhanced usual care rather than nothing, so the finding is not that attention beats neglect. The setting is routine primary care rather than a research clinic, which is the setting the result has to generalise to. The provider is a lay counsellor, which is the whole point of the design. And the two conditions are different, one a mood disorder and one a substance-use problem, which suggests the approach is a method rather than a treatment for one condition.

The proper use of these trials in an answer is as evidence about *delivery*. They establish that non-specialists, trained and supervised, can deliver a manualised psychological treatment in an Indian primary-care setting to a standard that outperforms the enhanced usual care they were compared against. They do not establish that any untrained worker can deliver anything, and they say

nothing about what happens to such a programme once the trial's training and supervision structure is withdrawn.

IN INDIA

Because the delivery question has been answered here rather than borrowed, the decision facing an Indian service is not whether non-specialist delivery can work but whether the district can supply the supervision that made it work. That converts the planning question into an arithmetic one a resident can actually do: how many workers, at what supervision frequency, in groups of what size, adds up to how many hours of somebody's week, and whose week is it. A programme that cannot answer that has not been designed, and it will reproduce the trials' structure for as long as its initial enthusiasm lasts.

One further caution belongs with the Indian evidence and is the kind of point that separates a good answer from a fluent one. Trials of this sort recruit, train and supervise their workers to a standard that a research team can guarantee and a district usually cannot, so the trial estimates what the model can do under favourable conditions rather than what it will do under ordinary ones. That is not a criticism of the trials, which were conducted in routine settings precisely to narrow that gap. It is a reminder that the implementation question, whether the supervision structure survives once the researchers leave, is a separate empirical question from the efficacy question, and that this chapter's sources answer the first and not the second.

27.8 The design decisions, assembled

DECISION	THE QUESTION TO ANSWER	THE COMMON FAILURE	WHAT IT COSTS IF YOU GET IT RIGHT
Which task	Is it bounded, rule-governed, checkable, and escalatable	Redistributing judgment rather than a procedure	Time spent writing the boundary down before anybody is trained
Which worker	What competencies are required, and who can acquire them	Selecting on qualification rather than on demonstrable competency	A selection process, and the willingness to reject candidates
What training	Short, repeated, case-based, with observed practice	A single workshop, cascaded through tiers, measured by attendance	Repeated contact rather than one event
What supervision	Frequency, supervisor tier, group or individual, protected or voluntary	Dropped when a post falls vacant, and never restored	A standing claim on the time of the tier above the worker
What decision support	Presentation-based sequences with explicit exclusions and referral points	Distributing the document and calling it implementation	Adaptation work, and alignment with the local formulary
What adaptation	Change the surface, preserve the safety architecture	Preserving the materials and abandoning the supervision	Explicitly recording what was changed and why
What evaluation	Outcome in the person, not activity in the worker	Reporting numbers trained instead of people improved	An outcome measure collected routinely from the start

PITFALL

Counting workers trained as the programme's result. Trained workers are an input, and they decay: people leave, get redeployed, or stop doing the task when nobody asks about it. A programme reporting cumulative numbers trained across several years is reporting a quantity that has no relationship to how many people are delivering the task today. The number that matters is how many trained workers are currently active and supervised, and almost nobody publishes it.

MNEMONIC

Task, Trainee, Training, Tools, Talk-to-someone

The five things a task-sharing programme has to specify: the bounded task, the worker who will do it, the training that produces the competency, the decision-support tools that hold the reasoning, and the supervision and escalation route. The last is the one that disappears, and it is the one all the evidence depended on.

Bounded	Enhanced usual care	Routine primary care	Supervision
THE PROPERTY THAT MAKES A TASK REDISTRIBUTABLE	THE COMPARATOR IN BOTH INDIAN LAY-COUNSELLOR TRIALS	THE SETTING THOSE TRIALS RAN IN, NOT A RESEARCH CLINIC	THE COMPONENT CUT FIRST AND THE ONE THE EVIDENCE RESTED ON

Task sharing redistributes a bounded task together with its training, its decision support, its supervision and its escalation route, and the two Indian lay-counsellor trials that outperformed enhanced usual care in routine primary care are evidence about that whole package rather than about the counsellor alone.

28 · Telepsychiatry: models, telepsychiatry guidelines and Indian regulatory framework

Telepsychiatry earns its marks as a service, not as a gadget. The examinable question is never whether a psychiatrist can be seen on a screen; it is what a health system gains, and what it must build, when the specialist and the patient stop having to be in the same building. That reframing decides everything downstream: who is stationed where, who is permitted to start the encounter, what is counted, and what the counting is allowed to prove. Candidates lose the marks in three predictable places. They describe the modalities of remote consultation, which is a technology answer to a service question. They recite a national helpline total as though it were a number of patients. And they attribute a governance position to an instrument whose current status they have not checked. The professional-conduct frame around remote practice, the consent and confidentiality obligations at a distance, and the identity of the operative regulatory code all belong to the Forensic ethics, consent and legal procedure notes, and are not restated here. What follows is the delivery architecture, the national programme, the monitoring discipline and the equity argument.

28.1 The service problem, stated before any solution

A specialist mental-health workforce is small, unevenly distributed and slow to grow. Every scheme for extending it is a variation on the same manoeuvre: move the scarce expertise, or move the decision that needs it, so that the specialist's time is spent only on what nobody else can do. Referral moves the patient. Outreach moves the specialist. **Task-sharing** moves the task. **Remote consultation** moves the encounter itself, and it is the only one of these that does not consume travel time at either end.

That is the whole economic case, and it explains why the design questions are the ones they are. If the constraint is specialist time, a remote service that simply converts face-to-face appointments into video appointments has bought convenience and no coverage: the same clinician sees the same number of people, and the queue behind them is unchanged. Coverage only rises when the remote link lets someone who is *not* a psychiatrist hold the patient contact while the psychiatrist supplies the decision. A telepsychiatry programme that has not answered who holds the contact has not yet been designed.

RULE **A telepsychiatry programme is judged by the population it can now reach and by whose time it saves, never by the technology it deploys.** Bandwidth, platform and picture quality are preconditions. They are not the intervention, and an appraisal that stops at them has appraised nothing that could change a coverage figure.

28.2 Hub and spoke, and what the architecture is actually doing

The dominant service design for remote specialist care is the **hub-and-spoke model**. A small number of **hubs** concentrate the scarce resource, in this case psychiatrists and the supervisory and record-keeping apparatus around them. A much larger number of **spokes** sit where the population is, staffed by whoever the health system already has at that level. The link between them carries clinical decisions in one direction and clinical information in the other.

Read structurally, the model does three separable things, and a programme can succeed at one while failing at the others. It *concentrates* expertise, so that a psychiatrist is not the rate-limiting step in every district simultaneously. It *distributes* access, so that the distance a patient travels falls to the nearest spoke rather than the nearest specialist. And it *standardises*, because a hub that supervises many spokes tends to impose one assessment format, one prescribing habit and one escalation route across all of them. The third effect is the one most often overlooked and is frequently the largest clinical gain, because it raises the floor of care at the periphery rather than the ceiling at the centre.

The corresponding failure modes are equally separable. A hub with too many spokes becomes a new bottleneck with a longer waiting list. A spoke with no trained staff becomes a video booth, which returns the model to moving the patient. And a link with no agreed **escalation route**, meaning a pre-arranged answer to who attends an acute presentation at the spoke and where the person goes, leaves the periphery holding risk it was never resourced to hold. That is where remote services generate their most serious incidents.

This chapter cannot state the official Indian architecture in the terms its own programme documents use. What was sought was a directly retrieved official passage defining the hub-and-spoke or tiered service structure for Indian telepsychiatry, including the roles of peripheral cells and specialist hubs; the programme and institutional sources approached either failed to open, exceeded retrieval limits or returned server errors, and no such passage was obtained. The reasoning above is service-design reasoning and is offered as such. An answer that describes the generic architecture and then says plainly that the official Indian tiering is not established from a source it has read is a stronger answer than one that supplies a confident tier list.

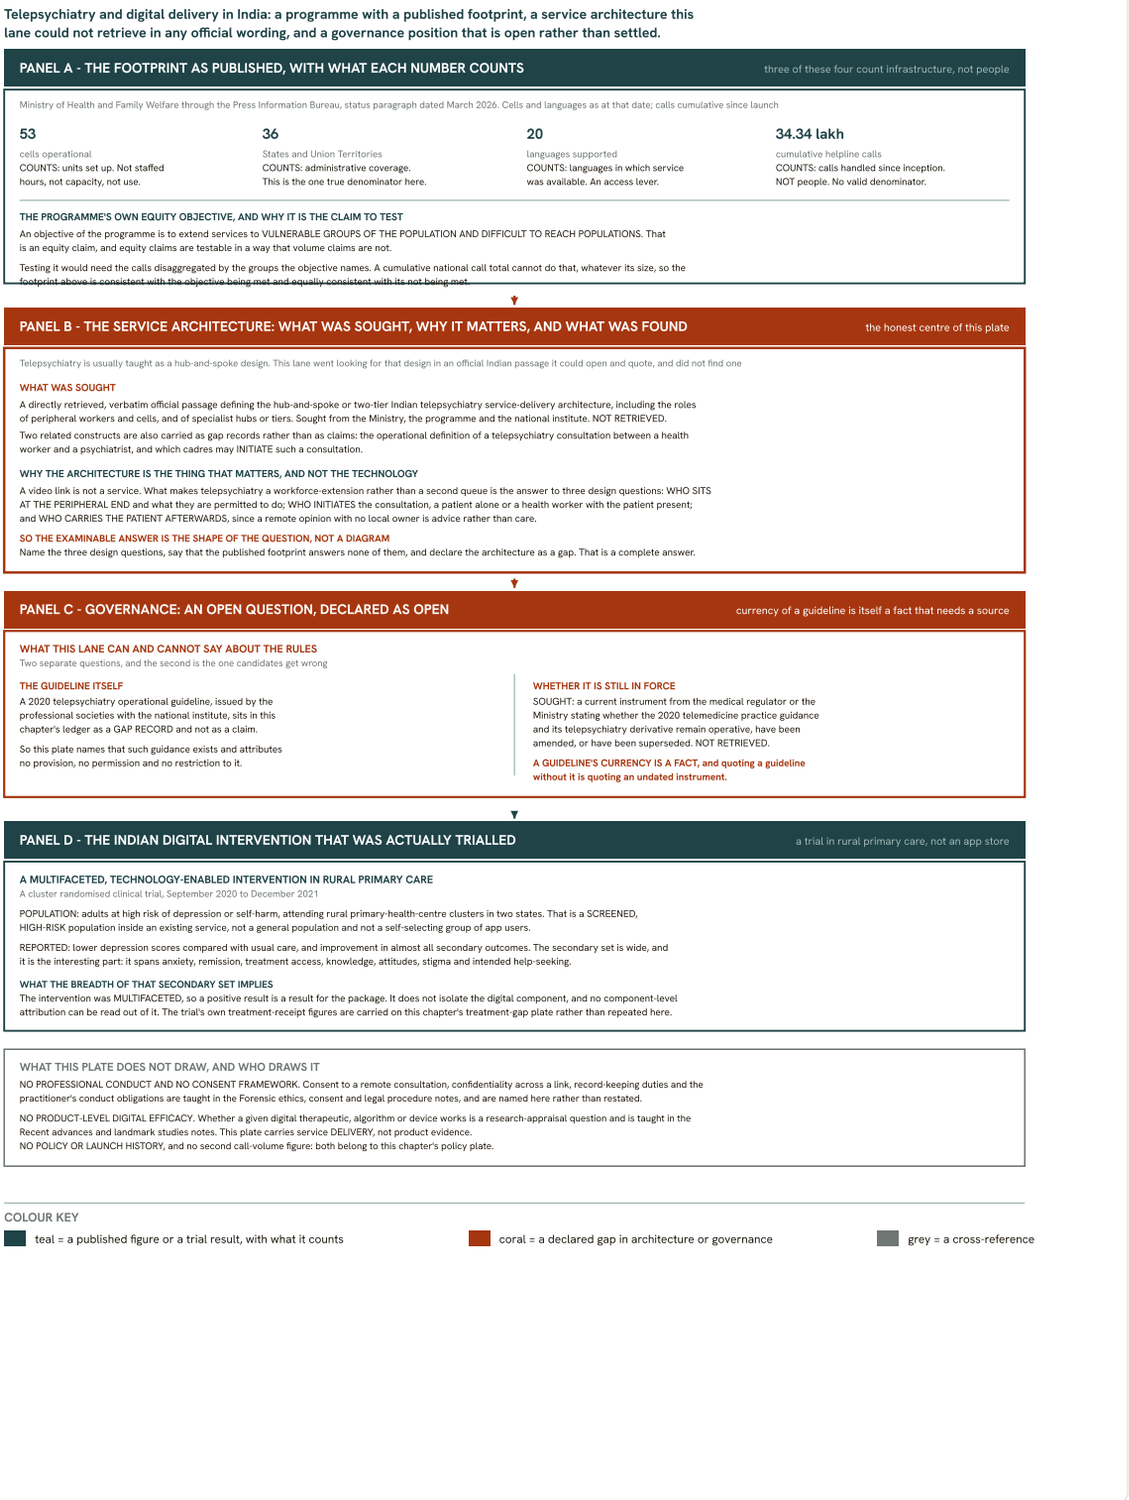


Fig 39.10 Telepsychiatry and digital delivery: a published footprint over an architecture that is a declared gap. The programme reports 53 cells across 36 States and Union Territories with support in 20 languages, and more than 34.34 lakh cumulative helpline calls; three of those four count infrastructure and the fourth counts calls rather than people. Its stated objective is to extend services to vulnerable and difficult-to-reach populations, which the footprint can neither confirm nor refute because it is not disaggregated. The service architecture itself is the honest centre of the plate: a verbatim official passage defining the hub-and-spoke or two-tier design, and the roles of peripheral workers and specialist hubs, was sought and not retrieved, and no current instrument confirming whether the 2020 telemedicine guidance remains operative was retrieved either. The one Indian trial drawn here reported lower depression scores and improvement in almost all secondary outcomes in a high-risk rural primary-care population, for a multifaceted package that does not isolate its digital component.

MNEMONIC

Place, Person, Permission, Proof

The four axes on which any remote mental-health programme is specified, and the order in which they fail. Where are the spokes and where is the hub. Who staffs each end. Who is permitted to initiate the encounter. What is measured well enough to show whether the reach was real.

28.3 The national instance, and what it was set up to do

India's national tele-mental-health service is delivered through **Tele-MANAS**, and its own statement of purpose is the right place to begin, because a programme's declared objectives are what its monitoring should be reporting against. Among those objectives is the explicit commitment to extend services to vulnerable groups of the population and to populations that are difficult to reach. That sentence is worth reading as a design instruction rather than as a preamble: it makes equity of reach a stated aim of the programme, which means underserved reach is a legitimate thing to hold the programme to, not an external criticism of it.

The programme's published national status figures describe its footprint. At its most recent published national status update, reported by the Ministry of Health and Family Welfare, there were **53** operational cells across **36** States and Union Territories, with service available in **20** languages, and **more than 34.34 lakh** calls received and addressed through the helpline since the service began.

53	36	20	more than 34.34 lakh
OPERATIONAL CELLS	STATES AND UNION TERRITORIES COVERED	LANGUAGES OF SERVICE	CALLS ADDRESSED SINCE INCEPTION, NOT PEOPLE

Each of those quantities counts a different thing, and the differences are the examinable content. A cell is a staffed unit, so the cell count is a supply figure. Coverage of States and Union Territories is an administrative reach figure and says nothing about depth within any one of them. The language count is an access figure and is arguably the most under-discussed of the four, because a helpline in a language the caller does not think in is a helpline the caller does not use. The call total is a **cumulative activity figure** and is the one that is misread most often, because activity is not the same object as the **unique users** a coverage claim would need.

28.4 Reading a programme's own numbers without laundering them

A national figure is only a fact when four things travel with it. Who issued it, what period it refers to, which population it describes, and what it is a proportion of. Strip any one and the number becomes a rumour with a decimal point in it. Applied to a helpline total, the fourth question is fatal: a cumulative count of calls has no valid denominator at all. It is not divided by the population, because it is not a count of people; it is not divided by callers, because one caller may call many times; and it cannot be compared with a prevalence, because the two count incommensurable things.

WHAT THE FIGURE IS	WHAT IT CAN SUPPORT	WHAT IT CANNOT SUPPORT
Number of operational cells	A statement about staffed service capacity	Any claim about how much of that capacity was used
States and Union Territories covered	Administrative reach across the country	Depth of reach inside any one of them
Languages of service	A concrete access provision for callers	That every caller reached service in their own language
Cumulative calls addressed	Total activity handled since the service opened	Number of individuals helped, or any rate per population

TRAP **Reading a cumulative helpline call total as a count of people served, and then dividing it by a population to produce a coverage rate.** Candidates do this because the number is large, official and easy to remember, and because every other figure in a public-health answer is a rate. The correct sentence names the count as calls received and addressed since the service began and stops there. If an examiner wants a coverage claim, the honest answer is that the programme publishes activity, not unique users, so a coverage rate cannot be computed from it.

IN INDIA

Because reaching vulnerable and difficult-to-reach populations is a stated objective of the national tele-mental-health programme, a state cell planning its rota has a specific and answerable design question rather than a general aspiration: which of the populations named in that objective is it currently failing to receive, and what in the service is producing the failure. Language of service, hours of operation and the referral destination offered at the end of the call are the three levers that can be moved locally, and the objective makes each of them a monitoring question rather than a courtesy.

28.5 The regulatory frame, and the currency question inside it

Remote practice in India operates inside a national telemedicine framework and a professional-conduct code, and this chapter names them as the governing frame and goes no further. The provisions themselves, the regulation numbers, the sequence by which the newer conduct regulations were notified and their subsequent status, and the identity of the code that is operative today are taught in the Forensic ethics, consent and legal procedure notes, together with consent to the modality, identity verification at both ends, confidentiality at a distance and the situations in which a remote encounter is not enough. Repeating any of that here would produce a second and possibly divergent account of a statutory position, which is the single most damaging thing a service chapter can do to a legal one.

Two regulatory facts a service chapter does need, and cannot supply, are declared here rather than filled in from recollection. The first is the identity, issuing body and year of the dedicated Indian telepsychiatry operational guidance that practitioners commonly refer to; what was sought was that document itself, and it was not established from a source this chapter opened. The second is more consequential and is a genuine live uncertainty: no current national instrument was retrieved that states explicitly whether the standing telemedicine practice

guidance and the telepsychiatry guidance derived from it remain operative, have been amended, or have been superseded. Multiple official routes were attempted and each failed to open or returned unusable content.

The teaching point survives the gap, and is arguably sharpened by it. Guidance for remote practice has changed repeatedly and quickly in every jurisdiction that has issued any, so the currency of the instrument is part of the clinical question and not an administrative footnote. A candidate who names the framework, states that its current standing must be checked against the issuing body rather than assumed, and points to the chapter that owns the conduct code has answered correctly. A candidate who recites a version and a year from memory has asserted the one thing most likely to have moved.

28.6 Who may open the encounter, and why that answer sets the ceiling on coverage

Return to the architecture. The hub extends the workforce only if a non-specialist may hold the patient and initiate the contact with the specialist. Everything about the size of a programme therefore depends on a permission: whether a nurse, an auxiliary nurse midwife, an allied health professional or another cadre of health worker may start a telepsychiatry consultation, and under what supervision. If only a registered medical practitioner may initiate, the spokes must be staffed by doctors and the model inherits the shortage it was built to relieve. If a trained health worker may initiate, the spoke can sit at the level where the population actually presents.

The specific Indian operational position on this is a declared gap in this chapter. Two related propositions were sought and not established from an opened source: an operational definition of telepsychiatry service delivery conducted between a health worker and a psychiatrist, and an operational statement permitting nurses, allied health professionals, auxiliary nurse midwives and other health workers to initiate such a consultation. Both are exactly the kind of proposition a candidate is likely to half-remember from a slide, and both change the shape of a service completely depending on which way they read.

What can be taught safely, and what an examiner is really testing, is the consequence structure. The permission determines the staffing of the spoke. The staffing determines the training burden. The training burden determines how fast the programme can scale and how uniform its quality will be. And the same permission determines the medico-legal question of who is responsible for the patient during and after the link, which is where the conduct code becomes unavoidable and where the reader is sent to the chapter that owns it.

28.7 Governance, monitoring and the evaluation a remote service owes

Remote delivery makes some things easier to measure and some things easier to lose. It is easier to count contacts, because they pass through a system that logs them. It is harder to know what happened afterwards, because the person who would have noticed is at the other end of a link that has closed. Monitoring has to be designed against that asymmetry rather than around the data that happen to accumulate.

A remote mental-health service should be able to state the following about itself, and a service that cannot is being appraised on its activity log alone.

- Who was reached, described by the populations the programme says it exists to reach, and not only by total volume.
- What proportion of contacts ended with a named destination, meaning a specific onward service, review or clinician, rather than advice alone.
- Whether the destination was reached, which requires a follow-up mechanism the link itself does not provide.
- How acute risk identified during a remote contact was escalated, to whom, and with what agreed local response.
- What the spokes needed and did not get, which is the only signal that reliably distinguishes an under-trained spoke from an under-used one.

The Indian reach argument for remote specialist care, including the contribution of the non-governmental sector to extending it, is developed in the Schizophrenia: management, antipsychotics and adverse effects notes, and is not re-derived here.

PITFALL

Designing the escalation route last. A remote service is commissioned around the consultation, and the acute pathway is added afterwards if at all, so the spoke discovers during its first emergency that nobody has agreed who attends, where the person is taken, or which clinician holds responsibility until they arrive. The escalation route is not an appendix to a telepsychiatry protocol; it is the part that determines whether the service can safely accept unfiltered presentations at all.

28.8 Equity, and the way a remote service quietly reproduces the gap it was built to close

Remote delivery is often presented as an equity intervention, and it can be one, but only against an explicit account of who is excluded. Every remote service imposes a set of entry requirements that face-to-face care does not: a device, a connection, sufficient privacy at the patient's end to speak freely, enough literacy or confidence to navigate a call, and a language the service offers. Each of those requirements is unevenly distributed in the same direction as the underlying disadvantage, so an unexamined remote service can expand access fastest among those who already had the most.

This is why the programme's own equity objective matters operationally. If reaching vulnerable and difficult-to-reach populations is a declared aim, then the entry requirements above become the programme's business, and the counter-measures are concrete: service in more languages, a staffed spoke where the person has no private device, hours that fit shift and agricultural work, and a call destination that exists in the caller's own district. None of that follows automatically from putting a service online, and the difference between a remote service that narrows a gap and one that widens it lies almost entirely in whether those counter-measures were built.

Telepsychiatry extends a workforce only when someone other than the specialist may hold the patient; everything else about the model, including its numbers, follows from that permission.

29 · Collaborative and stepped-care models of mental health delivery

Most people with a common mental disorder are already in contact with a health service; almost none of them are in contact with a psychiatrist. Collaborative and stepped care are the design answers to that arithmetic. They are not treatments, and they are not referral pathways with better manners. They are ways of organising a primary-care system so that the decisions requiring specialist judgement reach a specialist while the contact, the follow-up and most of the delivery stay where the patient already is. The marks here sit on structure rather than on content: what the roles are, what triggers a change of treatment, where the joins are, and what the system must measure to know whether any of it worked. Case management and assertive community treatment as service models, with their caseload logic and target populations, belong to the Schizophrenia: management, antipsychotics and adverse effects notes; the general-hospital liaison models and the classic economic argument for funding a liaison service belong to the Consultation-liaison, geriatric and transcultural psychiatry notes, which also records that the Indian cost-effectiveness evidence for such services is what is missing. Neither is restated here.

29.1 Stepped care as a rationing principle, stated honestly

Stepped care organises an intervention set by intensity and offers the least intensive intervention that is likely to work, with an explicit rule for moving up when it does not. Two commitments make it work and both are frequently dropped in practice. The first is that the lowest step must be a real intervention with a real chance of benefit, not a holding position. The second is that the rule for stepping up must be written down before the patient arrives, applied on a schedule, and triggered by an observation rather than by the patient's persistence or the clinician's patience.

It is worth being candid that stepped care is a rationing principle. A system adopts it because it cannot give everybody the most intensive option, and it distributes the scarcest resource by escalating only those who have already failed something cheaper. That is a defensible allocation, and it is defensible precisely to the extent that the escalation actually happens. Where the stepping rule is absent, stepped care degenerates into a queue with a therapeutic vocabulary, and the people who do worst are the ones least able to keep coming back to ask.

A distinction worth fixing early is between the lowest step and no step at all. Watchful waiting, in which a clinician defers treatment and reviews later, is a legitimate clinical decision but it is not the bottom rung of a stepped-care system unless the review is scheduled and somebody owns it. The bottom rung of a working system is an active low-intensity intervention with a named deliverer, a defined content and a defined end point. The difference is not semantic: in the first arrangement the patient returns if they deteriorate, and in the second the service comes back to look.

There is also a stratification logic that stepped care is sometimes confused with. Stratified allocation assigns the patient at the outset to the intensity their presentation appears to warrant, and does not require failure at a lower level first. Stepped care allocates by observed response rather than by initial appearance. Most real programmes are hybrids, and it is worth saying so rather than pretending purity: severity, risk and previous treatment history are used to enter people at a higher step, while response governs movement thereafter. What must not happen is that the hybrid becomes an excuse to hold everybody at the bottom.

RULE **A stepped-care system is defined by its stepping rule, not by its steps.** Any service can list interventions in order of intensity. What makes the arrangement stepped care is a pre-specified observation, taken at a pre-specified interval, that obliges the system to change what it is doing. Without that, the arrangement is a menu.

29.2 Collaborative care, and the three roles that constitute it

Collaborative care is a specific organisational form and not a synonym for good communication. Its defining feature is that a named person, distinct from both the treating primary-care clinician and the specialist, takes responsibility for keeping a defined group of patients under active review. That role is usually called the **care manager**, and the model stands or falls on whether it exists as a funded post rather than as an aspiration attached to somebody's existing job.

Three roles interlock. The primary-care clinician retains clinical responsibility, prescribes and remains the patient's medical home. The care manager holds the register, makes the scheduled contacts, delivers or supports the low-intensity intervention, records the agreed measure and raises the cases that are not improving. The consulting specialist reviews those cases, usually without seeing most of the patients, and advises on the ones that are stuck. The specialist's time is therefore spent on decisions rather than on encounters, which is the whole point: one psychiatrist supervising a caseload through a care manager can influence a population many times larger than the one they could see.

Two structural features are easy to miss and are frequently examined. The first is the **population register**: collaborative care is delivered to an enumerated list of patients, not to whoever returns. A patient who stops attending is a case the register flags, whereas in ordinary care they are simply absent. The second is that the specialist review is **caseload-based** rather than referral-based. The specialist looks systematically at those who are not improving, which inverts the usual selection, where the cases that reach a specialist are the ones somebody remembered to refer.

MNEMONIC

Register, Role, Review, Route

The spine of any collaborative-care arrangement. An enumerated population register. A named care-manager role that is somebody's actual job. A scheduled review carrying an agreed measure. And a route that obliges the system to do something different when the measure does not move.

29.3 Measurement-based escalation, and what a measure is for here

The measurement in measurement-based care is not for diagnosis and not for research. It exists to make deterioration and non-response visible to a system that would otherwise notice neither. That purpose sets its properties: it must be short enough to be repeated at every contact, stable enough that a change means something, and administered by whoever is actually in the room, which in this model is the care manager and not the specialist.

The clinically important discipline is that the measure must be attached in advance to a consequence. A rating repeated at intervals with no pre-agreed action attached is documentation, and documentation absorbs the same clinical time as a decision while producing none. The service therefore has to specify, before the programme opens, what change counts as adequate progress, at which review the question is asked, and what happens on each answer. Naming and describing an instrument is legitimate teaching; this chapter typesets no instrument's items, response options, scoring rule or thresholds anywhere, and a candidate should describe the escalation logic rather than reproduce a cut-off.

Note where this places the burden. The system, not the patient, owns the trigger. In ordinary care the patient who is not improving has to come back and say so, which selects for confidence, proximity and money. In a measurement-based system, the review happens on schedule and the absence of improvement declares itself. That single inversion is where most of the equity gain in these models comes from.

Enumerated, not attending	Scheduled, not requested	Caseload, not referral	Obliged, not offered
WHO THE REGISTER COUNTS	WHEN THE REVIEW HAPPENS	HOW THE SPECIALIST SELECTS CASES	WHAT THE MEASURE TRIGGERS

Every tile in that strip names a default that ordinary care sets the other way, which is a useful way to audit an existing service: ask which of the defaults it has actually inverted.

29.4 The joins, which is where the system leaks

Steps are easy to design and joins are hard to run, so the joins are where an examiner will probe. Each transition between steps is a moment at which responsibility can be dropped rather than handed over, and the failure is usually silent because both sides believe the other is holding the patient.

THE JOIN	WHAT MUST BE AGREED BEFORE A PATIENT REACHES IT	HOW IT FAILS IN PRACTICE
Recognition into the programme	Who may enter a patient on the register, and on what basis	Entry depends on individual clinician interest, so the register reflects enthusiasm rather than need
Low-intensity step to structured intervention	The observation that triggers it and who acts on it	The trigger exists but no capacity was commissioned behind it, so the step is nominal
Structured intervention to specialist decision	Whether the specialist reviews the record or must see the patient	The review defaults to an appointment, and the model reverts to ordinary referral
Any step back down, or discharge	Who resumes routine follow-up and with what safety-netting	Discharge is treated as an ending, and relapse re-enters the system as a new presentation

The number of steps is not standardised and varies between programmes; what is invariant is the structure of the join. An answer that describes a named ladder of steps has described one programme. An answer that describes what must be agreed at each transition has described the model.

29.5 The Indian evidence, and what it licenses

The most directly relevant Indian evidence for this construct comes from public primary-care facilities in Goa, where recovery from common mental disorders improved among the patients attending them when collaborative care was delivered by **a trained lay counsellor**. The finding matters for a specific reason that is easy to state and easy to overstate: the care-manager role was filled by a trained lay counsellor rather than by a mental-health professional, in the public sector, in ordinary facilities.

Read carefully, that result licenses a workforce conclusion and not a universal one. Its population is patients with common mental disorders attending public primary-care facilities, so it says nothing about severe mental illness, about facilities without the supervisory structure the trial supplied, or about a counsellor cadre without the training the trial delivered. The delivery agent is the variable of interest, and the supervision is the condition that makes the delivery agent safe. A programme that copies the cadre and drops the supervision has copied the cheap half.

IN INDIA

Where a district cannot place a mental-health professional in every facility, and most cannot, the practical decision this evidence supports is to build the care-manager role from a trained non-specialist cadre with specialist supervision attached, rather than to hold the model in abeyance until specialist posts are sanctioned. The gain was obtained in ordinary public primary-care facilities with a lay counsellor delivering the collaborative intervention, so the design question for a district programme is what training and what supervisory link it can actually sustain, not whether a psychiatrist can be found for every centre.

29.6 What these models are not, and the errors that follow

Collaborative care is not consultation. In consultation the specialist answers a question about a patient and the arrangement ends there. In collaborative care the specialist joins a standing system with a register, a schedule and a named person holding the caseload, and the relationship persists whether or not any particular question is asked. Liaison practice and the taxonomy of liaison service models belong to the Consultation-liaison, geriatric and transcultural psychiatry notes.

Neither model is case management, and the distinction is examinable. Case management and assertive community treatment are service models for a different population, defined by their caseload logic and their target group, and they are taught with their figures in the Schizophrenia: management, antipsychotics and adverse effects notes. Case management appears here only as a component: the care-manager function inside collaborative care borrows the idea of a named person holding a defined caseload, and borrows nothing else.

TRAP **Describing stepped care as "start with the least intensive treatment", full stop.** Candidates reproduce the first half of the definition because it is the memorable half, and lose the marks that sit on the second. The examinable content is the escalation: what is measured, when it is measured, what threshold obliges a change, and who is responsible for making it. A description without the stepping rule is a description of a waiting list.

PITFALL

Commissioning the register and the care manager but not the step above. A programme that can identify and monitor many more people than the specialist tier can absorb generates a queue of correctly identified, correctly measured, non-improving patients with nowhere to go, and the care manager becomes the person who apologises to them. Capacity at the higher step is a precondition for opening the lower one, not an enhancement to be added later.

29.7 Fidelity, and why a model that worked in a trial stops working as a programme

These are complex interventions, and complex interventions degrade in characteristic ways when they scale. A trial supplies training, supervision, protected time and a research infrastructure that measures whether the protocol was followed. A programme supplies none of those unless somebody commissions them, and the components that are dropped first are always the invisible ones. Supervision is invisible. The register is invisible until it is missing. The scheduled review is visible only to the person who did not receive it.

The practical consequence is that a programme evaluation that reports outcomes without reporting **fidelity** cannot interpret its own result. If outcomes are poor and fidelity was poor, the model was not tested. If outcomes are poor and fidelity was good, the model failed in that setting, which is a genuinely different finding with a genuinely different response. Reporting both is not methodological fussiness; it is the difference between fixing the implementation and abandoning the design.

Three components carry most of the fidelity burden in these models and are worth naming as the things to protect when a budget is cut. The first is supervision of the non-specialist deliverer, because it is what converts a trained cadre into a safe one. The second is the scheduled review, because dropping it converts the register into a list. The third is the capacity at the step above, because escalation that has nowhere to go teaches the whole system to stop escalating. A programme that keeps the counselling sessions and loses all three has kept the visible half of the intervention and discarded the mechanism.

29.8 Evaluation, and the quantities the model has to be able to state

These models are unusually easy to evaluate honestly, because their own machinery generates the necessary denominators. The register defines who was in the programme, so proportions can be computed rather than estimated. A service running collaborative or stepped care should be able to state each of the following, and one that cannot is reporting activity.

- How many people were entered on the register in a stated period, and out of what attending population.
- What proportion received a scheduled review at the intended interval, which is the fidelity measure that matters most.
- What proportion of those not improving at review were actually stepped up, which is the single number that distinguishes the model from a menu.
- What proportion of specialist reviews were caseload-based rather than triggered by an individual referral.
- What happened to those discharged, since a model that improves recovery and loses relapse has moved the problem rather than solved it.

Costing and financing sit outside this construct and are taught in the mental health economics and financing material in these notes; the classic economic argument for funding a liaison service, and the observation that Indian cost-effectiveness evidence in this space is scarce, are already made in the Consultation-liaison, geriatric and transcultural psychiatry notes and are not re-derived or re-declared here.

> Collaborative care is defined by a register, a care manager and caseload-based specialist review; stepped care is defined by the rule that forces the step, not by the list of steps.

Special-setting and disaster mental health

30 · Disaster mental health: psychological first aid, phases and intervention models

A disaster does not produce one clinical problem; it produces a population with a changed distribution of need. Most of the people affected will recover without any formal intervention. A minority will develop diagnosable disorder. A separate and often forgotten group already had severe mental illness before the event and has just lost its supply of medicines, its clinic, its

records and its family carers. A group larger than any of these has been displaced, and **displacement** generates its own morbidity from crowding, lost livelihood and lost privacy rather than from the original event. A last group is the responders. Any response designed around only one of those groups fails all the others, and that is the examinable point: disaster mental health is a service-design question about matching a tiered response to a heterogeneous population, not a question about which therapy to give a survivor. The statutory machinery of disaster management, its tiers and its office-holders, and the professional and ethical frame around emergency work are taught in the Forensic ethics, consent and legal procedure notes and are not restated here. The clinical early-trauma intervention itself, with its action principles, its follow-up rule and its standing cautions about early single-session and pharmacological approaches, is taught in the Anxiety, OCD and trauma-related disorders notes. What follows is the population response and the service that has to deliver it.

30.1 Needs assessment as a method, not as an impression

The first deliverable of a mental-health response is not an intervention, it is a description of who needs what. **Rapid needs assessment** is a method with a defined output, and a response that skips it delivers whatever the responding agency happens to be able to provide, which is a common and expensive failure.

A usable assessment establishes several things that cannot be inferred from the event itself. It establishes the size and composition of the affected population, including how many are displaced and where they now are. It establishes what has been lost from the pre-existing service: which facilities are standing, which staff are contactable, which drug supply lines still run, and whether patient records survived. It establishes what the population itself says it needs, which in the acute period is reliably practical rather than psychological. And it establishes what already exists locally, meaning the community structures, faith networks and lay helpers who will still be present long after any external team has left.

The output of that assessment should be a distribution rather than a headline. A single figure for the proportion of a population needing mental-health care is almost always useless in operations, because it merges people who need information and practical help with people who need medication continued and people who need specialist treatment. The three are met by different tiers with different workforces and different costs.

MNEMONIC

Assess, Tier, Train, Track

The sequence a disaster mental-health programme is built in, and the order in which it is usually cut. Assess before responding. Tier the response so each level of need meets the cheapest workforce that can safely meet it. Train and support the people who will still be there. Track what happened, because a response nobody measured teaches the next response nothing.

30.2 The tiered psychosocial response

The organising structure of every credible disaster mental-health response is a **layered response**, in which the widest layer serves everybody affected and each successive layer serves a smaller group with a more specialised workforce. Read from the base upwards, the layers are these.

The base layer is basic services and security delivered in a way that protects dignity: shelter, food, water, safety, and above all information, because uncertainty about missing relatives and about what happens next is one of the largest sources of distress in the acute period. This layer is not usually thought of as mental-health work and is the layer with the greatest population mental-health effect. Above it sits community and family support: restoring contact between separated people, reactivating community structures, supporting normal routines for children, and enabling collective mourning practices. Above that sits focused, non-specialised support delivered by trained and supervised workers who are not mental-health professionals, aimed at those whose distress is not settling. At the apex sits specialised mental-health care for the minority with severe or persisting disorder, and for those whose pre-existing illness has destabilised.

Three properties of the layered structure carry most of the teaching weight. It is not a triage sequence: the layers operate simultaneously, not one after another. It is inverted in cost and in volume, so the layer serving most people is cheapest and the layer serving fewest is dearest, which is exactly why an unplanned response over-invests at the top. And each layer requires a different workforce, so a plan that names the layers without naming who staffs each one has not been costed.

This chapter cannot attribute the layered framework to an issuing body from a source it has opened, and does not do so. The structure is taught here as service reasoning; a candidate who describes the layers and their workforce implications has the substance, and one who attaches a body and a year from recollection has added the one element most likely to be wrong.

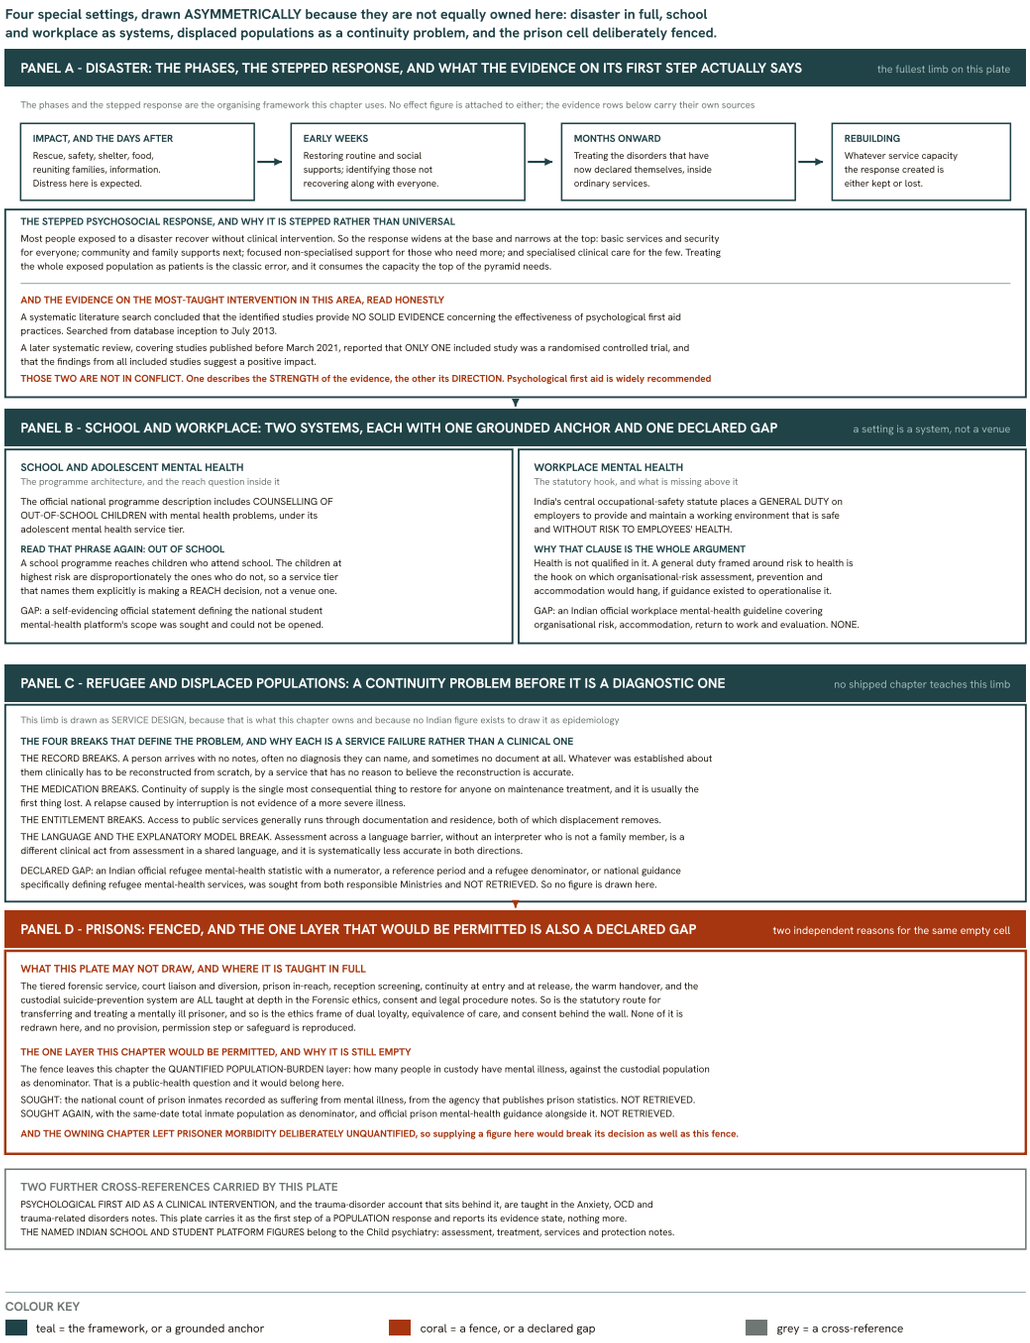


Fig 39.11 Four special settings, drawn to different depths because they are owned to different depths. Disaster carries the phases and the stepped psychosocial response in full, with its first step read honestly: one systematic search found no solid evidence of psychological first aid effectiveness while a later review found only one randomised trial among its included studies and reported that all of them suggest a positive impact, which are statements about strength and about direction rather than a contradiction. School rests on an official service tier that names counselling of out-of-school children, a reach decision rather than a venue one, with the national student platform's scope a declared gap. Workplace rests on a general statutory duty to maintain a working environment without risk to employees' health, with no Indian operational guideline above it. The displaced-population limb is drawn as four breaks in continuity, and no Indian refugee mental-health statistic was retrieved. Prisons are fenced to the Forensic ethics, consent and legal procedure notes, and the one permitted layer is itself a declared gap.

30.3 Where psychological first aid sits, and what it is for

Psychological first aid occupies the focused non-specialised layer. It is a structured, brief, humane response delivered by trained non-specialists to people in acute distress after a critical event, and its purpose is to reduce immediate distress, meet practical needs and connect the person with continuing support. Two boundaries define it more usefully than any description of its content.

The first boundary is that it is not a treatment. It does not target a disorder, it is not indicated by a diagnosis, and completing it is not a therapeutic episode. Its clinical action principles, the rule governing active monitoring and review afterwards, and the standing cautions about what must not be done in the acute phase are taught in the Anxiety, OCD and trauma-related disorders notes, and a candidate should point there rather than reproduce them.

The second boundary is that it is not universal. The people who benefit are those in acute distress, not everyone who was present. A response that pushes an emotional intervention onto every exposed person misallocates the workforce and risks pathologising a normal reaction, and it also crowds out the practical help that the acute period actually demands. The competence being tested when an examiner asks about psychological first aid is usually this discrimination, not the recall of a set of steps.

The delivery implication is the reason the model matters at scale. Because it is deliverable by trained non-specialists, it converts teachers, volunteers, community health workers and relief staff into a usable mental-health workforce for the widest tier of need. That is a workforce multiplication, and it is what makes a population-scale response arithmetically possible at all.

30.4 The evidence base, and how to hold it honestly

This is the part of the topic where candidates most often assert more than the literature supports, and it is worth getting exactly right, because the honest position is more impressive than the confident one.

A systematic search of the literature on this intervention concluded that nothing it identified amounted to solid evidence of effectiveness, which was the basis for its finding that guidelines could not be built on the evidence then available. A later systematic review of the same intervention reported a different but compatible picture: across what it included, randomized evidence was confined to a single trial while the studies as a whole pointed towards benefit.

Recommended widely	**Directionally positive**	**Randomized evidence scarce**	**No solid effectiveness evidence**
STATUS IN DISASTER RESPONSE PRACTICE	SIGNAL ACROSS INCLUDED STUDIES IN THE LATER REVIEW	DESIGN OF THE INCLUDED STUDIES	FINDING OF THE EARLIER SYSTEMATIC SEARCH

Those statements are not in conflict and reading them as a contradiction is the error to avoid. One reports the absence of solid effectiveness evidence; the other reports that the studies that do exist point in a favourable direction while almost none of them are randomized. Both are statements about the evidence base rather than about the intervention, and together they

describe a widely recommended intervention whose recommendation rests substantially on plausibility, consensus and the absence of harm rather than on trial evidence.

RULE **Absence of solid effectiveness evidence is not evidence of ineffectiveness, and it is not permission to assert effectiveness either.** The correct sentence names the intervention, states what the reviews found about the evidence, and gives the reason it is nonetheless recommended: it is low-risk, deliverable at scale by non-specialists, and the alternative on offer in the acute period is frequently nothing at all.

There is a structural reason the evidence is thin, and saying it demonstrates understanding rather than evasion. Randomizing an immediate humane response after a catastrophe is difficult to justify and difficult to execute; the exposed population is displaced and hard to follow; the outcome of interest emerges over a long period; and the intervention is delivered by many different people in many different forms under one name. Those are features of the setting, not failures of the researchers, and they mean the evidence base is unlikely to become deep quickly.

30.5 What changes as the response moves through its phases

A disaster response is not one service running for a long period; it is a sequence of different services with different staffing, and the transitions between them are where continuity is lost. The most important operational fact about the sequence is counter-intuitive: the mental-health workload does not peak when the media attention does. It rises as the acute relief effort demobilises, at precisely the point when external teams withdraw and funding attention moves elsewhere.

PERIOD OF THE RESPONSE	WHAT THE MENTAL-HEALTH SERVICE IS MAINLY DOING	THE CHARACTERISTIC FAILURE
Before any event, as preparedness	Building the plan, the roster and the training so that a response does not have to be invented	No plan exists, so the response is assembled by whoever arrives first
Immediate aftermath	Practical help, information, safety, and restoring supply and records for those already under treatment	Psychological intervention is offered while the practical layer is unmet
Relief and early recovery	Focused non-specialised support, case-finding for persisting distress, and support to responders	External teams deliver in parallel to the local service rather than through it
Later recovery	Specialist treatment for persisting disorder, and handover into the standing service	The external response has already left and the standing service was never strengthened

The service consequence of the last row is the single most useful thing to be able to say about disaster mental health. A response measured by what it delivered during the emergency will look successful and leave nothing behind. A response measured by what the district can do afterwards is a different programme, staffed and budgeted differently from the start.

The first row is the one under a service's own control and the one most often empty. **Preparedness** is the only phase in which decisions can be taken calmly, and it is the phase in which a district can decide, in advance, who leads the mental-health component, which staff are on the roster, what training they will have had, how a drug supply is re-established, where displaced patients will be seen, and what the mutual-aid arrangement with neighbouring districts is. None of those questions gets a good answer in the week after an event, and every one of them determines what the response is able to do.

The transition that matters most is the last one, and it has a name worth using precisely. **Handover** here does not mean transferring a list of patients; it means transferring a functioning arrangement, including who now holds each patient, where the record sits, who supervises the trained non-specialists, and which budget line carries the work from that point. A handover that consists of a summary document is an ending described as a transition.

> **PITFALL**
>
> The parallel service. An external team arrives with its own staff, its own protocol, its own records and its own funding line, works alongside the district service without integrating with it, and departs when the funding ends. Every patient it was holding is then either lost or handed to a service that has never seen them, has no record of the treatment given and was not resourced for the extra load. Working through the existing service is slower at every point and is the only arrangement that leaves capacity behind.

30.6 The responders are an exposed population

The workforce delivering the response is itself part of the affected population, and in a local disaster it is often part of it twice over. Rescue workers, health staff, police, volunteers and administrators are exposed to the event, to its aftermath and to sustained contact with distress, and in a district-level disaster many of them have lost homes and relatives in the same event they are working through.

Responder care is therefore a programme component and not a gesture, and it is designed at the level of the system rather than the individual. The measures that make a difference are structural: rostering that enforces rest rather than relying on individuals to request it; briefing before deployment about what will be seen; a named supervisor with an explicit responsibility for the team's condition; a mechanism to stand somebody down without it being treated as a failure; and access to care that is confidential enough that a local worker will actually use it in a place where everybody knows everybody.

The confidentiality point deserves emphasis because it is where responder programmes most often fail in practice. A support service staffed by the worker's own colleagues, in the worker's own institution, is a service the worker may not use for exactly the problems it exists to address. Designing an access route that sits outside the immediate hierarchy is a service-design decision that costs almost nothing and determines whether the provision is real.

30.7 Special groups the tiered plan has to name explicitly

A layered plan distributes generic need well and handles specific vulnerability badly unless the vulnerable groups are named in the plan itself. Each of the following requires a different arrangement, and each is routinely discovered late.

- People already receiving treatment for severe mental illness, whose immediate need is supply and continuity rather than a new intervention, and whose records may no longer exist.
- Separated and unaccompanied children, for whom reunification and protection are the mental-health intervention and everything else is secondary.
- Older and disabled people, whose access to a distribution point, a shelter or a clinic may be the binding constraint on every other service they need.
- The bereaved, particularly where bodies were not recovered, since the practical and administrative aftermath of an unconfirmed death is itself a prolonged stressor.
- Displaced people living in temporary settlements, where crowding, loss of privacy and loss of livelihood generate new problems that the original event did not.

Naming these groups in the plan does more than allocate resources. It commits somebody to counting them, and a group nobody counts is a group whose needs will be met by whatever is left over.

TRAP **Answering a disaster mental-health question as though the task were to treat post-traumatic stress disorder in survivors.** Candidates go straight to a diagnosis and a therapy, and thereby miss the population structure, the layered response, the people whose established treatment has been interrupted, and the responders. The disorder-level question is a legitimate one and belongs to the chapter that owns acute trauma; the question asked here is what a service does for a population, and the marks are distributed accordingly.

30.8 Evaluation, and what a response should be able to show afterwards

Disaster responses are evaluated less often than almost any other public-health intervention, for understandable reasons and with predictable consequences: each response is improvised, and the improvisation is not recorded, so the next one starts from the same place. Building the evaluation into the plan before the event is the only point at which it is cheap.

What a response should be able to state afterwards is modest and rarely available. How many people were reached at each layer, and out of what estimated affected population. What proportion of those already under treatment for severe mental illness had their treatment continued without interruption, which is the most concrete continuity measure available and the one most likely to reflect real service performance. How many non-specialists were trained, how many were supervised afterwards, and how many were still in post once the response ended. What proportion of responders used the support offered. And what capacity remained in the district when the response closed, compared with what was there before it opened.

That last quantity is the honest measure of a disaster mental-health programme, and it is the one that distinguishes a relief operation from a health-system intervention. An event that leaves a

district with more trained workers, a written plan and a functioning referral route has produced a durable gain. An event that leaves a district with a completed activity report has not.

> The response is layered and simultaneous, psychological first aid is one rung of it delivered by trained non-specialists, and the population that most reliably needs a psychiatrist is the one that already had one before the disaster.

31 · School mental health programmes and adolescent mental health services

The school is the only place where a health system can reach almost an entire age group without anyone having to decide they are unwell. That single property is what makes school mental health a public-health topic rather than a paediatric one, and it is what the marks here are actually about: the school is a delivery platform, and the examinable questions are the ones asked of any platform. Which tier of need is being addressed. Who delivers at each tier. What proportion of the intended population is genuinely reached. Who pays, and for how long. And how anybody would know whether it worked. What a school programme contains, how life-skills teaching is built, how a single session becomes a school culture, and how an identified child is assessed, treated, safeguarded and referred are taught at length in the Child psychiatry: assessment, treatment, services and protection notes, which also carries the Indian delivery platforms with their coverage figures. None of that is reproduced here. What follows is the platform read as a system: tiers, coverage, financing, evaluation, and the adolescents the school platform structurally cannot reach.

31.1 The school as a platform, and what the platform property buys

Every population intervention faces a selection problem. Services delivered in clinics reach the people who present, and presentation is filtered by recognition, permission, money and distance, all of which are unevenly distributed. A **platform intervention** avoids the filter by delivering where the population already is. Schools are the most complete such platform available for young people, and their advantages are structural rather than clinical.

The assembled population is the first advantage: no health-seeking decision is required, so the intervention reaches those who would never have been brought. The second is repeated contact over a long period, which is what any preventive intervention needs and what a clinic almost never has. The third is the presence of a trained adult workforce that already observes the child daily and is capable, with support, of noticing change. The fourth is that the school is itself part of the causal environment, so it can be modified rather than merely used as a venue.

The disadvantages are equally structural and are less often stated. The platform reaches only the enrolled and the attending, so it systematically misses the young people at highest risk. Teachers are not a free workforce and every task added to them displaces another. The school's institutional interest is educational attainment, so a mental-health programme is retained when it is seen to serve that interest and dropped when it is not. And a programme delivered inside a school inherits the school's own hierarchies, so confidentiality is harder to guarantee than in a clinic, which is a matter for the chapter that owns the identified child.

RULE **A school mental-health programme is a platform decision before it is a content decision.** The question that determines its effect is not which curriculum is delivered but which tier of need it is aimed at, who delivers it, and what proportion of the intended age group is actually enrolled, attending and present when it is delivered.

31.2 The prevention tiers applied to a school

The prevention taxonomy is developed as a framework in the prevention material of these notes, and its definitions belong elsewhere again; what belongs here is its application, because the school is the setting in which all its tiers can be seen operating in one building. Applied to a school population, the tiers separate cleanly by whom they are aimed at.

TIER	WHO IT IS AIMED AT IN A SCHOOL	WHO DELIVERS IT, AND WHAT IT COSTS
Universal	Every pupil, irrespective of risk, and the school environment itself	Teachers within the timetable; cheapest per head and dearest in total, because everyone receives it
Selective	Groups with a raised exposure but no identified problem, defined by circumstance rather than by symptom	Trained school staff or a visiting worker; requires a defensible basis for grouping and careful handling of labelling
Indicated	Individual pupils showing early problems that fall short of a disorder	A trained non-specialist with supervision, or a visiting mental-health worker; smallest numbers, highest unit cost
Treatment, which is not prevention	Pupils with an established disorder	The health service, not the school; the school's job here is recognition and a working referral route

The last row is where most school programmes fail, and it is a system failure rather than a content failure. A universal programme in a school raises recognition, and raised recognition generates identified children. If no referral destination exists, the programme has manufactured demand it cannot serve, and the teachers who identified the children carry the consequence. Building the highest tier before opening the widest one is the single most consistent lesson from school programmes in any health system.

MNEMONIC

Enrolled, Attending, Present, Reached

The four filters standing between an intended school population and the population a programme actually serves. Coverage claimed against enrolment is an administrative figure. Coverage claimed against those present when the intervention was delivered is a service figure. The gap between them is where school programme evaluations most often go wrong.

31.3 The Indian delivery platforms, named and located

India does not lack school and adolescent platforms; it has several, they sit in different ministries, and their coverage provisions, training structures and target ages are taught with their figures in the Child psychiatry: assessment, treatment, services and protection notes. This chapter names them so that a reader knows what exists and where the detail lives, and reproduces none of their coverage figures.

- The district mental health programme's school outreach component, which places mental-health work inside schools through the district team and its master-trainer model.
- Rashtriya Kishor Swasthya Karyakram, the national adolescent health programme, which carries adolescent mental health inside a wider adolescent health remit and includes a clinic-based adolescent-friendly service tier.
- The Ayushman Bharat School Health and Wellness Programme, which works through designated teachers delivering scheduled health-promotion sessions within the school timetable.

The system-level observation worth making about that list is the one a candidate can make without any figure at all. The platforms are plural, they sit under different ministries with different reporting lines, and they were designed at different times for different primary purposes. A district therefore does not have one adolescent mental-health programme with a coverage figure; it has several partial provisions whose overlaps and gaps nobody is required to reconcile. Integration across those platforms is the practical Indian problem, and it is an administrative problem rather than a clinical one.

31.4 The adolescents who are not in school

Every advantage of the school platform is conditional on enrolment, and the conditionality is not random. Young people out of school are disproportionately those who are working, married early, disabled, displaced, from the poorest households, or already in difficulty. On any plausible account of risk they are the group with the highest need, and the platform designed to reach adolescents reaches them least.

This is not an oversight in the Indian service description. The **National Mental Health Programme**'s own account of what its district teams deliver on outreach includes the counselling of children who are not in school and who have mental health problems, listed within its adolescent mental health work. The provision is therefore already named at national level, which changes the nature of the question a district faces: it is not whether an out-of-school limb is legitimate, but what contact point it uses.

That contact point is the whole design problem, because there is no assembled population to work with. The available options are all second-best and each carries a different reach and a different cost. Work through the health facility, which returns the selection filter the school platform existed to avoid. Work through the community health worker, which reaches households but adds to an already loaded role. Work through workplaces where adolescents are employed, which reaches the working young but requires an employer's cooperation. Work through

community and religious institutions, which reaches those inside them. Work through a remote or helpline route, which requires a device and a private moment. A district plan that names one of these and resources it has an out-of-school limb; a plan that names none has a school programme and a claim.

IN INDIA

Because the national programme's own outreach description already includes counselling for children outside school who have mental health problems, the practical decision for a district team is not whether to serve this group but which single contact point it will commit to and staff. Choosing one route and running it properly reaches more of this population than distributing the intention across all of them, and it makes the coverage question answerable, because a named route has a denominator that can be counted while a general intention does not.

31.5 Coverage, and the difference between a programme existing and a programme reaching

Coverage is the quantity on which any platform intervention is properly judged, and it is the quantity most often replaced by a proxy. A programme can be sanctioned in a district and absent in its schools. It can be present in a school and delivered to a fraction of classes. It can be delivered to every class and reach only those present on the day. And it can reach everyone present and change nothing, because the tier addressed was not the tier of need.

The honest formulation asks four questions in sequence and refuses to collapse them. What is the intended population, defined by age and by geography rather than by enrolment. What proportion of that population is enrolled and attending, which is an education-system figure the health system must obtain rather than assume. What proportion of those actually received the intervention as designed. And what proportion of those identified as needing the next tier reached it. A programme that can answer only the first two is reporting the education system's performance, not its own.

Districts	Schools	Teachers trained	Adolescents reached
WHAT A SANCTION FIGURE COUNTS	WHAT AN INSTITUTIONAL FIGURE COUNTS	WHAT A TRAINING FIGURE COUNTS	THE ONLY DENOMINATOR THAT ANSWERS COVERAGE

Read that strip from left to right as a descent from what is easy to report towards what a population actually received. Each tile counts a real thing, and only the last one has the population in its denominator. A programme report that stops before the last tile has not answered the coverage question; it has answered an administrative one and left the reader to assume the rest.

TRAP **Quoting the number of schools or districts covered as though it were population coverage.** Districts covered is an administrative reach figure, and its denominator is districts, not children. Schools covered is likewise a figure about institutions. Neither can be converted into a proportion of the adolescent population without the enrolment and attendance data that turn one into the other, and the young people at highest risk sit precisely in the gap between those denominators.

31.6 What an adolescent service tier has to look like to be used

An adolescent service is not a paediatric service with older patients or an adult service with younger ones, and the reasons are about access rather than about diagnosis. Adolescents present late, present through somebody else, and stop attending for reasons that have more to do with the shape of the service than with their symptoms. Designing an **adolescent-friendly service** means removing the specific obstacles that this age group faces, and each of them is a system decision.

Location and hours are the first, because a service open only during school or working hours obliges the young person to explain their absence, which is itself a disclosure. The entry route is the second: a service reached only through a parent excludes the adolescents whose difficulty involves the household, and a service reached only through a psychiatric department attaches a label at the door. Staff attitude is the third and is the one adolescents themselves report as decisive. And a clear statement of what will and will not be shared with a parent or a school is the fourth, since an adolescent who cannot predict where their words will go will manage the risk by saying less. The confidentiality boundary, the limits of the teacher's role and the handling of an identified child are taught in the Child psychiatry: assessment, treatment, services and protection notes and are not restated here.

The last design problem is **transition**. A young person who reaches an adolescent service and remains unwell must eventually move into adult services, and that move is a documented point of disengagement across health systems. Treating transition as an administrative transfer rather than as a planned handover with an overlap period, a named receiving clinician and a shared plan produces a predictable loss of exactly the patients who most needed continuity. A district that builds an adolescent tier without deciding what happens at its upper edge has built a service with an exit and no destination.

31.7 Financing, and why school programmes have a characteristic life cycle

School mental-health programmes fail in a recognisable pattern, and the pattern is financial rather than clinical. A programme is introduced with project or scheme funding that pays for training and for a coordinating post. The training is delivered and counted. The coordinating post ends with the project. The trained teachers are transferred, promoted or retire, and the school reverts. A subsequent evaluation finds no effect and concludes that the intervention does not work, when what happened is that it was never sustained long enough to be tested.

The financing questions that determine whether this happens are worth naming precisely, because they are answerable in advance. Whether the recurring costs, meaning refresher training and supervision, sit on a recurring budget line or a project one. Whether the coordinating role belongs to a post or to a person. Whether the receiving health service at the top tier was funded to take the referrals the programme generates. And whether the education system carries any of the cost, since a programme entirely funded by the health side is a programme the school has no institutional stake in retaining.

PITFALL

Counting teachers trained as the programme's output. Training is an input, it is easy to count, and it is what appears in reports. It tells you nothing about whether anything was delivered afterwards, whether the trained teacher is still in that school, or whether a single child reached a service. A programme reporting only training numbers has reported its expenditure.

31.8 Evaluation, and what this chapter cannot establish

The evaluation of a school mental-health programme should be aimed at the tier it addresses, and mismatched evaluation is the commonest reason a reasonable programme is judged ineffective. A universal intervention should be judged on population-level change and on the school environment, not on the recovery of identified children, of whom it will contain few. An indicated intervention should be judged on the outcome of those identified, and on whether they reached the next tier. Judging a universal programme by clinical outcomes in a handful of identified pupils guarantees a null result whatever the programme did.

Two things this section would ordinarily state are declared as gaps rather than supplied, and both are exactly the kind of fact that feels certain from memory.

The first is the national student mental-health support initiative run from the education side. What was sought was a primary official statement defining its scope for school, college and university students; the official routes approached timed out or returned no retrievable page text, and no such statement was obtained. The construct exists and is named here; its scope, its year and its provisions are not established by anything this chapter has read, and a candidate is better served by saying so than by attributing a remembered description to a ministry.

The second is prevalence. No defensible Indian figure is available here for cyberbullying among school-going children and adolescents, and none is supplied. This matters as method and not only as caution: a prevalence figure measured in one population cannot be transferred to another simply because both consist of school children, and a figure quoted without its survey, its year, its age band and its definition of the behaviour is not a fact about India. The correct answer names the construct, says the Indian figure is not established, and moves to what a service should do about a problem whose size it does not know.

A school programme's effect is set by its tier and its coverage, and its coverage is capped by enrolment, which is precisely why the out-of-school limb is not an afterthought but the part that reaches the highest risk.

32 · Workplace and occupational mental health

Work is where most adults spend most of their waking, non-domestic life, and it is the one setting in which an institution can change a person's exposure by changing its own arrangements. That is why workplace mental health is a public-health topic and not an occupational-medicine footnote: the employer holds levers no clinician holds. It can change how much is demanded, how much control a person has over it, whether support exists, whether the

role is intelligible, and whether the job will still exist next month. A clinician can treat the consequences of all of these and can change none of them. The examinable content sits on that asymmetry, and the commonest failure in an answer is to describe workplace mental health as a set of services offered to employees, which relocates the problem into the worker and leaves the work untouched. What follows is the organisational risk, the prevention structure applied to a workplace, the accommodation and return-to-work system, the manager's role and its limits, and what an Indian programme can and cannot be held to.

32.1 The setting, and the three populations inside it

A workplace, viewed as a health platform, has the same assembled-population advantage as a school and an additional one: the institution has a continuing relationship with the person, a record of their attendance, and a financial interest in their capacity to work. It also has three distinguishable populations, and a programme that does not separate them will be judged against the wrong outcome.

The first population is everybody employed, well and unwell, whose exposure is set by how the work is designed. The second is those with emerging difficulty who are still working, often at reduced effectiveness and usually undisclosed. The third is those absent from work or returning from absence, whose problem is no longer only clinical but also contractual, financial and social. Each requires a different intervention, delivered by a different part of the organisation, and evaluated on a different measure.

Two further groups sit at the edges and are routinely forgotten. Workers in insecure, informal or contract arrangements are exposed to more of the hazards and covered by fewer of the provisions, so a programme built for permanent employees can widen a gap inside its own workforce. And managers, who are asked to carry much of the delivery, are themselves employees with the same exposures.

32.2 The statutory floor, and how far it reaches

India's central occupational-safety statute places a general duty on employers, which is to provide and to keep in place a working environment that is safe and does not put employees' health at risk. The wording of the duty is the load-bearing part, and it is worth reading twice. It concerns the *environment*, not the worker's constitution. It runs to *health*, without a qualifier confining it to physical health. And it obliges the employer to *provide and maintain*, which is a continuing obligation rather than a one-off provision.

A caution about how far that reading may be pressed, and it applies to any statutory duty. The exact strength of a duty, including whatever qualification the statute attaches to it, is the substance of the obligation and not decoration around it; a duty framed absolutely and a duty framed with a practicability qualifier are different claims about what an employer must do. This chapter states the duty and does not reproduce the statutory text, the section locator or the qualifying words, and a candidate arguing a legal point should go to the statute rather than to a summary of it. The disability entitlement frame that sits behind adjustments in employment is taught in the Mental health legislation and forensic psychiatry notes.

RULE **Workplace mental health is a work-design intervention first and a service offering second.** Interventions aimed only at the individual, however good, leave the exposure in place and locate the problem in the person exposed to it. An answer that lists counselling services without addressing how work is organised has described the smaller half of the field.

32.3 Organisational risk, described as properties of work rather than of workers

The **psychosocial hazards** of work are features of how a job is arranged, and they are describable independently of who holds the job. The categories that recur across occupational settings are these, and each has a corresponding organisational remedy that no individual intervention can substitute for.

- Demand: the volume, pace and difficulty of the work relative to the time and resources given for it.
- Control: how much say the worker has over how, when and in what order the work is done.
- Support: whether practical help and supervision are actually available from managers and colleagues.
- Role: whether what the person is responsible for is defined, consistent and free of conflicting instructions.
- Relationships: exposure to bullying, harassment, discrimination and unresolved conflict.
- Change and security: how organisational change is managed, and whether the job itself is precarious.

The important structural claim about these hazards is that they interact rather than add. High demand alongside high control is a demanding job; high demand alongside low control is the arrangement most consistently associated with harm. Effort without corresponding reward, recognition or security behaves similarly. Named theoretical models formalise these interactions and attach evidence to them; this chapter does not name them or their originators, because no claim row here carries them, and a model attributed from memory to the wrong author is exactly the kind of error a viva finds. The categories above are usable without the model names, and an answer that reasons from demand and control to a specific organisational remedy has demonstrated the understanding the model exists to convey.

MNEMONIC

Design, Detect, Adjust, Return

The four functions of a workplace mental-health programme, in the order of their population effect and the reverse of the order in which they are usually built. Design the work so it does less harm. Detect difficulty early and make disclosure survivable. Adjust the job so the person can keep doing it. And build a route back for those who have stopped.

32.4 The prevention tiers applied to a workplace

The prevention taxonomy is developed as a framework in the prevention material of these notes and is applied here rather than restated. Applied to a workplace it separates the interventions by

the level at which they act, and the separation is what makes an evaluation interpretable.

LEVEL AT WHICH IT ACTS	WHAT IT LOOKS LIKE IN A WORKPLACE	WHO HAS TO ACT	HOW IT CHARACTERISTICALLY FAILS
The work itself, before anyone is unwell	Workload and rostering decisions, clarified roles, managed change, an enforced anti-harassment position	Senior management, because these are budget and structure decisions	Delegated to a health or welfare function that cannot change any of them
The whole workforce, as capability	Awareness that reduces the cost of disclosure, and training that makes managers approachable	The organisation, with time released for it	Delivered as a single session and never refreshed, so the culture reverts
Individuals with emerging difficulty	Early conversation, confidential access to help, adjustment before absence occurs	Line managers, supported by an occupational health or counselling route	Access exists but disclosure is unsafe, so nobody uses it until they are absent
Those already absent or impaired	Structured return to work, graded duties, sustained adjustment, continuity with treatment	Manager, occupational health and the treating clinician together	Handled as an attendance problem by a process designed for discipline

The job	The person	Disclosure cost	Return route
WHAT AN ORGANISATIONAL INTERVENTION CHANGES	WHAT AN INDIVIDUAL INTERVENTION CHANGES	WHAT DETERMINES WHETHER EITHER IS USED	WHAT DECIDES IF ABSENCE BECOMES EXIT

32.5 Adjustment, and why it is cheaper than everything downstream

Adjustment means changing the job so that a person with a mental-health difficulty can continue to do it, and its most useful property is that it is usually inexpensive and reversible. The adjustments that matter in mental health are unglamorous and mostly concern time, predictability and load: altered start and finish times so that treatment appointments or medication effects are accommodated; a temporary reduction in the most demanding component of the role rather than in the role as a whole; a quieter or more predictable working location; a phased increase in duties; a named person to check in; and a temporary change in the frequency of supervision so that difficulty is caught before it accumulates.

Two design principles carry most of the effect. The first is that adjustment should be available before absence occurs, not only as a condition of return, because the cheapest point to intervene is while the person is still at work. The second is that an adjustment must have a review date and a person who owns it, or it either lapses silently or becomes permanent by default, and both outcomes damage the arrangement's credibility for the next person who needs one.

IN INDIA

Because the employer's general duty concerns a working environment that is safe and without risk to employee health, and that duty is not written as a duty about physical health alone, the practical decision for an Indian organisation is to locate mental-health provision inside the occupational safety and health management system it already runs rather than beside it as a welfare activity. That placement is not cosmetic: it puts the provision where the statutory duty, the reporting line and the standing safety committee already sit, which is what converts an initiative into an obligation somebody is answerable for.

32.6 Return to work, treated as a system rather than a date

Return to work is where workplace mental health becomes measurable and where most organisations perform worst, because absence is managed by processes designed for a different purpose. The default machinery is **attendance management**, whose logic is verification and escalation, and applying it to a mental-health absence reliably delays return and sometimes prevents it.

A functioning return system has a small number of features and they are all about contact and structure. Contact during absence is maintained by a named person on an agreed basis, since silence is read as abandonment and repeated formal letters are read as threat. A return is planned before it happens, with the returning person present at the planning. Duties are **graded** rather than restored in full on the first day. The adjustments are written down, given a review point, and communicated to whoever needs to know and to nobody else. And the return is reviewed after a period rather than treated as complete on arrival, because relapse in the first weeks back is common and is best met with a further adjustment rather than a further absence.

PITFALL

Requiring full fitness before return. The demand looks prudent and is counterproductive: it converts a person who could do part of the job now into a person doing none of it indefinitely, removes the structure and contact that support recovery, and lengthens absence to the point at which returning at all becomes unlikely. A **capacity-based return**, which asks what the person can do and adjusts around what they cannot, is both safer and faster.

32.7 The manager's role, its limits, and the confidentiality problem

Most workplace mental-health programmes rest on line managers, and most manager training fails by being too ambitious about what a manager is for. A manager is not being trained to assess, diagnose, counsel or treat, and training that implies otherwise produces either an anxious manager who avoids the conversation or a confident one who does something inappropriate.

The manager's job has three parts and stops there. Notice a change in the person's work or presentation. Have a conversation that makes disclosure safe and is about work rather than about diagnosis. And connect the person to a route, whether that is occupational health, a counselling service or their own doctor, while adjusting what can be adjusted meanwhile. Training aimed at

those three does measurably more than training aimed at recognition of disorders, and it is also the version a manager can practise without exceeding their competence.

Confidentiality is the structural problem that determines whether any of this is used. An employee discloses to someone who also decides their appraisal, their duties and possibly their continued employment, and the same information that unlocks an adjustment can be read as evidence of unreliability. The resulting **cost of disclosure**, meaning what the employee reasonably expects to lose by speaking, is the single variable that determines whether any provision in the organisation is used at all. The design answer is to separate what the manager needs to know from what they must not: a manager needs to know what adjustment is required and for how long, and does not need a diagnosis, a treatment or a prognosis. Where an **occupational health function** exists, its correct output to a manager is a statement of capacity and adjustment, not a clinical account. Where a service is delivered by the employer's own staff, the same problem arises inside the service, which is why an access route independent of the reporting line is a design requirement rather than a courtesy.

TRAP **Reporting uptake of an employee counselling service as the outcome of a workplace mental-health programme.** Uptake is an activity figure with no population denominator attached, and it can rise for good reasons or bad ones: a workforce under increasing strain generates more uptake, as does a service that has become easier to reach. It says nothing about exposure, about absence, about return, or about the people who did not disclose. The outcomes that answer the question are absence and return-to-work measures, the proportion of adjustments requested that were granted, and any measure of the psychosocial hazards themselves.

32.8 Evaluation, and what is not established for India

The evaluation problem here is that the easily available measures belong to human resources and the meaningful ones do not exist unless somebody builds them. Absence data are collected everywhere and are a poor single measure, because absence falls both when people get better and when they become too insecure to take leave. A defensible evaluation therefore reads several quantities together: absence and its duration, the proportion of absences that end in return rather than exit, the proportion of requested adjustments granted and sustained to their review point, and a repeated measure of the psychosocial hazards themselves, which is the only quantity that reports on the intervention that matters most.

Two things a section on Indian workplace mental health would ordinarily supply are declared here as gaps. There is no primary-official Indian workplace mental-health guideline available to this chapter addressing organisational risk, prevention, reasonable accommodation, return to work, manager training and programme evaluation. What was sought was exactly such an instrument from the labour and health ministries and their technical bodies; the official sources approached either carried only a bare mention of workplace stress management without any framework, or established a general health duty without an express mental-health framework, or failed to open. The construct is named here without a national instrument behind it, and a candidate who says that plainly is answering correctly.

There is likewise no national coverage figure available: how many Indian workplaces run any mental-health provision, what proportion of the workforce it reaches, and what proportion of that workforce is in the informal arrangements that no such programme touches are all unquantified here. The absence is itself instructive, because a workforce whose majority is informally employed cannot be covered by an employer-based model however good that model is, and a national workplace mental-health strategy that addresses only formal employment addresses a minority of the working population.

> The employer's duty runs to the working environment and to health without qualification, so the first intervention in workplace mental health is always to the design of the work, and every service offered to the worker is downstream of it.

33 · Prison, refugee and other special-population mental health

A special population, in service terms, is not a population with an unusual diagnosis; it is a population whose access to care is controlled by something other than its need. For prisoners that something is custody. For refugees and displaced people it is legal status, documentation and mobility. For others it is the absence of an address, a guardian or an entitlement. The clinical material is largely the same material taught everywhere else in psychiatry; what changes, and what carries the marks, is that the ordinary route to care has been interrupted by an institution or a legal category, so the service has to be redesigned around the interruption. This section teaches the displaced-population limb in full, and for prisons only the population and system questions: what a prison mental-health figure would have to be to mean anything, and how prison mental-health systems differ across countries. Custodial law, prisoner rights, the ethics of practice behind the wall, the design of an in-reach service, continuity at entry and at release, and the custodial suicide-prevention system are taught at depth in the Forensic ethics, consent and legal procedure notes, and none of them is restated here.

33.1 What makes a population special, stated as a rule about access

The useful generalisation is that in each of these populations a third party stands between the person and the service, and that third party's interests are not identical to the patient's. A prison authority is accountable for security. An immigration authority is accountable for status determination. A shelter or an institution is accountable to whoever funds it. None of them is primarily accountable for the person's mental health, and each can, without malice, produce a decision that removes the person from care.

Three consequences follow and they recur across every population in this section. Continuity is fragile, because the person can be moved, released, transferred or deported without reference to their treatment. Consent and confidentiality operate under constraint, because the person's circumstances limit what a refusal can mean and who else has an interest in the information. And ordinary measures of service performance mislead, because the denominator is a population the health service does not control and often cannot count.

RULE **Where a third party controls access, service design is mostly about the joins with that third party and only secondarily about the clinical content.** The question that determines whether a special-population service works is not what treatment it offers but what happens to a patient at the moment the custodian, the authority or the institution makes a decision about them.

33.2 Displacement, read as a sequence of exposures rather than an event

Refugee and displaced-population mental health is often taught as though the trauma were the flight. That framing is incomplete and, for service design, misleading. The exposures are distributed across a sequence, and the parts a receiving service can actually change lie at the end of it.

Before departure sit the conditions that produced the displacement: conflict, persecution, disaster, destitution, and the losses accumulated while deciding to leave. During movement sit the hazards of the journey itself, separation from family, and, for many, exploitation and violence en route. After arrival sit a different set of exposures that are chronic rather than acute: uncertain legal status, prohibition or absence of work, dependence, crowded or temporary accommodation, discrimination, and the sustained absence of information about what will happen next.

The post-arrival exposures are the ones a health service and a policy can modify, and they behave differently from the pre-arrival ones. The events before and during flight are fixed and historical. The conditions after arrival are ongoing and administrative, and prolonged uncertainty about status functions as a continuing stressor rather than a background condition. That distinction is why an intervention aimed only at past trauma, in a person whose present circumstances remain unresolved, so often disappoints: it treats the part of the sequence that has ended and leaves the part that is still happening.

Two further groups belong in this limb and are frequently excluded from it. **Internally displaced people**, who have not crossed a border, are usually more numerous than refugees and are covered by fewer specific provisions. And people whose displacement is old, meaning long-settled populations in **protracted situations**, have a different problem again: their acute needs have passed and their structural disadvantage has become permanent.

33.3 What a service for displaced people actually has to solve

The design problems are concrete and are mostly not clinical. Each of them has a service answer, and a plan that has not answered them will deliver care to whoever happens to arrive able to use it.

- Entitlement: what a person is allowed to receive depends on a legal status that may be undetermined, contested or absent, so the service must decide in advance what it will provide to somebody whose status is unresolved.
- Language: assessment through an untrained **interpreter**, or through a family member, changes what is said and what is recorded, and interpreting capacity is the single most common binding constraint on the quality of a refugee mental-health assessment.

- Documentation: people arrive without records, so treatment history, previous diagnoses and current medication have to be reconstructed, and continuing supply of an existing medicine is frequently the most urgent clinical task.
- Mobility: populations move between camps, cities and countries, so a plan requiring several appointments at one location may be unusable, and a **portable record** held by the patient is worth more than a good record held by the clinic.
- Calibration: distress is expected and is not by itself a disorder, so a service must be able to distinguish the person who needs housing, information and a resolved status from the person who needs treatment, without dismissing the second or medicalising the first.

MNEMONIC

Status, Speech, Supply, Setting

The four things to establish before any assessment of a displaced person is worth recording. What legal status they hold and what it entitles them to. In what language the encounter will actually happen, and through whom. What treatment was already running and whether it can be continued. And what their living conditions now are, since those are the exposure a service can still change.

PITFALL

Building a separate refugee mental-health clinic funded by a project. It is visible, it is fundable, and it produces a parallel service with its own records, staff and closing date, attached to a population that is mobile and whose entitlement to the mainstream service was never resolved. When the project ends, the patients belong to nobody. Strengthening the general service to receive this population is slower, less visible and the only version that survives.

33.4 What is not established about refugee mental health in India

This is a section where a confident figure would be easy to produce and impossible to defend, so the position is stated plainly. No Indian primary-official refugee mental-health statistic is available to this chapter, meaning no figure carrying a numerator, a reference period and a refugee denominator; and no national guidance specifically defining refugee mental-health services was retrieved either. What was sought was exactly that, from the home and health ministries. The official sources reached described camps, relief schemes and legal procedure, and contained no mental-health statistic and no mental-health-specific guidance for this population.

The absence is worth teaching as content rather than passed over, because it explains a real feature of the field. India hosts displaced populations of several origins under several different administrative arrangements, and a person's entitlement can depend on which arrangement they fall under rather than on what they need. Where no single national mental-health provision defines the service, what a displaced person receives is determined locally, by whichever combination of the general health service, a state arrangement and a non-governmental provider happens to exist where they are. A candidate who can say that, and can say that the national figures do not exist, is describing the system accurately. A candidate who quotes a prevalence for

refugee mental disorder in India is quoting something no Indian official source in this chapter's reach has published.

Not established	Not retrieved	Not retrieved	Not available
INDIAN REFUGEE MENTAL-HEALTH STATISTIC WITH A REFUGEE DENOMINATOR	NATIONAL GUIDANCE DEFINING REFUGEE MENTAL-HEALTH SERVICES	NATIONAL COUNT OF INMATES RECORDED AS MENTALLY ILL	SAME-DATE INMATE POPULATION AS A DENOMINATOR

33.5 Prisons, and the boundary of what belongs here

Prison mental health is taught as a service in the Forensic ethics, consent and legal procedure notes, and the boundary is worth stating so a reader does not go looking for it here. That chapter carries the custodial service architecture and the standards it is held to. The statutory route by which a mentally ill prisoner is transferred and treated is named here only as the dependency it is, since a service whose transfer route does not work has no upper tier, and the reader is sent to the Forensic ethics, consent and legal procedure notes for the provision itself. The rights and ethics frame governing care behind the wall is likewise named here only as the standard against which any prison mental-health system is measured, and is taught in the same place.

What remains for a public-health chapter is narrower and genuinely distinct: the question of what a prison mental-health figure would have to be to mean anything, and the comparison of prison mental-health systems across countries. Both are population and system questions rather than service-design questions, and neither is taught elsewhere in these notes.

33.6 The numbers problem, and what an honest prison figure would require

It is well established, and stated in the chapter that owns the custodial service, that psychiatric morbidity in custody runs far above the community rate. Nothing in this section contradicts that direction. What this section does is explain why attaching a number to it is harder than it looks, and that explanation is examinable in its own right.

Start with what is routinely published. Prison administrations publish **administrative counts**, and a count of inmates recorded as mentally ill is a count of recognition and recording, not a count of illness. It depends on who did the recognising, with what training, at what point in the sentence, using what threshold, and on whether recording it had any consequence for the institution. A system with no reception assessment records few cases; a system with an active assessment records many; and the difference between the numbers describes the two systems, not the two populations.

Then note the denominator problem, which is specific and fatal to casual comparison. A prison population is a **flow** as much as a **stock**: the number held on a given date is not the number who passed through in a year, and the two differ most for the remand population, which is also the population with the highest turnover and often the highest morbidity. A numerator collected across a period, divided by a population counted on one date, produces a figure that is not a proportion of anything.

An Indian figure that would satisfy this chapter's standard therefore needs a numerator with a stated case definition and ascertainment method, a same-date or same-period denominator from the same source, and a stated reference period. That figure was sought here from the national prison statistics published under the home ministry by the National Crime Records Bureau, together with the corresponding total inmate population. It was not obtained: the national publication timed out on repeated attempts, and where an official route did serve, it exceeded the retrieval limit so that no verbatim table and row could be quoted. A related route supplied a total inmate population without the mental-illness numerator that would have to sit above it. No figure is printed here, and the correct answer in a viva is that the direction is established, the Indian proportion is not, and the reason is a numerator and denominator problem rather than an absence of interest.

TRAP **Quoting an administrative count of prisoners recorded as mentally ill as though it were a prevalence.** The count is a recognition figure produced by the prison system's own procedures, its denominator is usually a single-date headcount while the numerator accrues over a period, and the case definition is rarely stated. Candidates reach for it because it is the only number available. The stronger answer names the count for what it is, states that morbidity in custody runs far above the community rate, and explains what would have to be true for a real proportion to be computed.

33.7 How prison mental-health systems differ across countries

Comparing systems rather than services is the other thing this section can legitimately do, and the comparison is more instructive than a table of prevalences would be, because the structural choices explain most of the variation in what is delivered.

AXIS ON WHICH SYSTEMS DIFFER	ONE ARRANGEMENT	THE OTHER ARRANGEMENT	WHAT THE CHOICE DETERMINES
Who is accountable for health inside the walls	The prison administration provides health care itself	The health ministry provides it, inside a facility run by another department	Whether clinical standards and staffing follow health-service norms or custodial ones
Who employs the clinicians	Employed by the custodial service	Employed by an external health provider working in the prison	Independence of clinical judgement, and how easily posts can be recruited to
Where the sickest are treated	Within dedicated units inside the custodial estate	By transfer out to health-service facilities	Capacity limits, waiting for transfer, and who carries the risk while waiting
How coverage is reported	Activity counts from the custodial system	Health-service reporting with a defined population	Whether coverage can be compared with the general population at all

The last axis is the one that makes cross-national comparison so difficult, and it is worth saying explicitly. Where prison health is reported through the custodial system, its figures are not

commensurable with the health system's, so international differences in reported prison morbidity partly measure differences in reporting architecture. A comparison that does not say which architecture each country uses has compared reporting systems and called it epidemiology.

33.8 The other populations, and the structure they share

The bank of special populations is longer than these two, and the shared structure introduced at the start is what makes them teachable together rather than as a list. In each case a third party or a legal category stands between the person and ordinary care, and the service response is a set of joins rather than a new treatment.

People without a home have no address, which is the hidden requirement in most systems for registration, appointment letters, follow-up and medicine supply. People in institutions of any kind have their access mediated by the institution's staff, so their care depends on somebody else recognising a need and acting on it. People who have been trafficked or are in exploitative work have their movement controlled by another person, so confidentiality and the ability to attend are both compromised. Migrant workers within a country lose entitlements attached to their place of origin without gaining those attached to their place of work. In every case, the correct question is the same: at which administrative moment does this person fall out of care, and what would have to exist for them not to.

That question generalises, and it is the reason this material belongs in a public-health chapter at all. Special-population psychiatry is not a set of exotic clinical entities; it is the ordinary service, examined at the points where its assumptions about address, identity, entitlement, language and freedom of movement quietly stop being true.

> In every special population the third party who controls access is the design problem, and the figures these populations generate describe the recognition system that produced them before they describe the population itself.

Prevention, economics and digital mental health

34 · Levels of prevention (primary, secondary, tertiary) applied to mental health

A prevention taxonomy is a routing device, and it earns its place only because it stops the same argument being made several times over. Community psychiatry is full of interventions that sound alike and act at completely different points: a public awareness campaign, a training programme for general doctors, a helpline, a restriction on access to a lethal method, a supported employment service. Placed on the prevention axes, each one declares whom it acts on and at what point in the course of illness it acts, and once that is declared the argument about it becomes tractable. That is the work this section does for the rest of the chapter: it states the mapping onto mental health and then routes each application to the topic that owns it, so that no application is explained twice. The definitions of the terms themselves, with their lineage and the caution against reading the axes as one, are taught in the Intellectual disability and autism

spectrum disorder notes. The machinery of screening, meaning the criteria a screening programme must satisfy and the biases that inflate its apparent benefit, belongs to the Research Methods, Biostatistics and Epidemiology notes. Neither is restated here.

34.1 What the taxonomy is for, and the two questions it asks

The prevention vocabulary exists because two entirely different questions can be asked of any preventive intervention, and answering only one of them produces confusion that looks like disagreement.

The first question is temporal: at what point in the natural history of the disorder does this intervention act? Before it exists, once it exists but before it has declared itself, or after it is established and the target is now disability rather than onset? That axis yields primary, secondary and tertiary prevention.

The second question is about the recipient: who is offered this intervention, and on what basis were they selected? Everybody in a population, a subgroup at raised risk, or an individual already showing early signs? That axis yields **universal**, **selective** and **indicated** prevention.

The axes are independent, and that independence is the entire reason both exist. An intervention has a position on each, and knowing one position tells you nothing reliable about the other. Reading them as a single ladder is a well-known error with its own exam trap, and both the definitions and that trap are taught in the Intellectual disability and autism spectrum disorder notes. What this section supplies is the mental-health mapping and the routing that follows from it.

MNEMONIC

Whom, When, What, Whether

The four questions that make a prevention claim assessable. Whom does it act on, and how were they selected. When in the course of the disorder does it act. What does it actually consist of, delivered by whom. And whether anybody measured the outcome it is supposed to change, in the population it was aimed at.

34.2 The temporal axis, mapped onto mental health

Mapping the temporal axis onto psychiatry is less tidy than mapping it onto infectious disease, and the untidiness is itself worth understanding rather than glossing.

Primary prevention in mental health acts on determinants and exposures before any disorder exists. In practice this means the material conditions and the social environment, most of which sit outside the health sector: household poverty, violence in the home, adverse experiences in childhood, alcohol availability, working conditions, and access to lethal means. The awkward feature is that almost none of these is a psychiatric intervention, and almost all of the effective levers belong to other ministries. That is not a reason to demote primary prevention; it is a reason to say plainly that a psychiatric service's primary-prevention role is largely advocacy, intersectoral work and the design of population-level restrictions, not clinical activity.

Secondary prevention acts once the disorder exists but before it has declared itself or done its damage, and its content in mental health is early detection followed by prompt effective treatment. This is where most of a health service's preventive effort actually sits, because the interval between onset and treatment in psychiatry is long and the harm accrues during it. Note the boundary carefully: the machinery of screening as a discipline, the criteria a screening programme must meet and the biases that make an unevaluated programme look successful, are taught in the Research Methods, Biostatistics and Epidemiology notes. The stage-by-stage design of a screening-to-treatment pathway, and the rule that a screen without a named destination is not a service, are worked end to end in the Perinatal/women's mental health and emergency psychiatry notes. What belongs here is only the placement of secondary prevention as a tier, and the mental-health examples of what is looked for and by whom.

Tertiary prevention acts after the disorder is established, and its target is disability, relapse and social exclusion rather than onset. In psychiatry this is the largest of the three limbs by volume of work, because the conditions that dominate a service are long-term and their principal harm is functional. Almost everything a rehabilitation service does sits here, and it is rarely called prevention, which is precisely why it loses arguments about resources to interventions that carry the word.

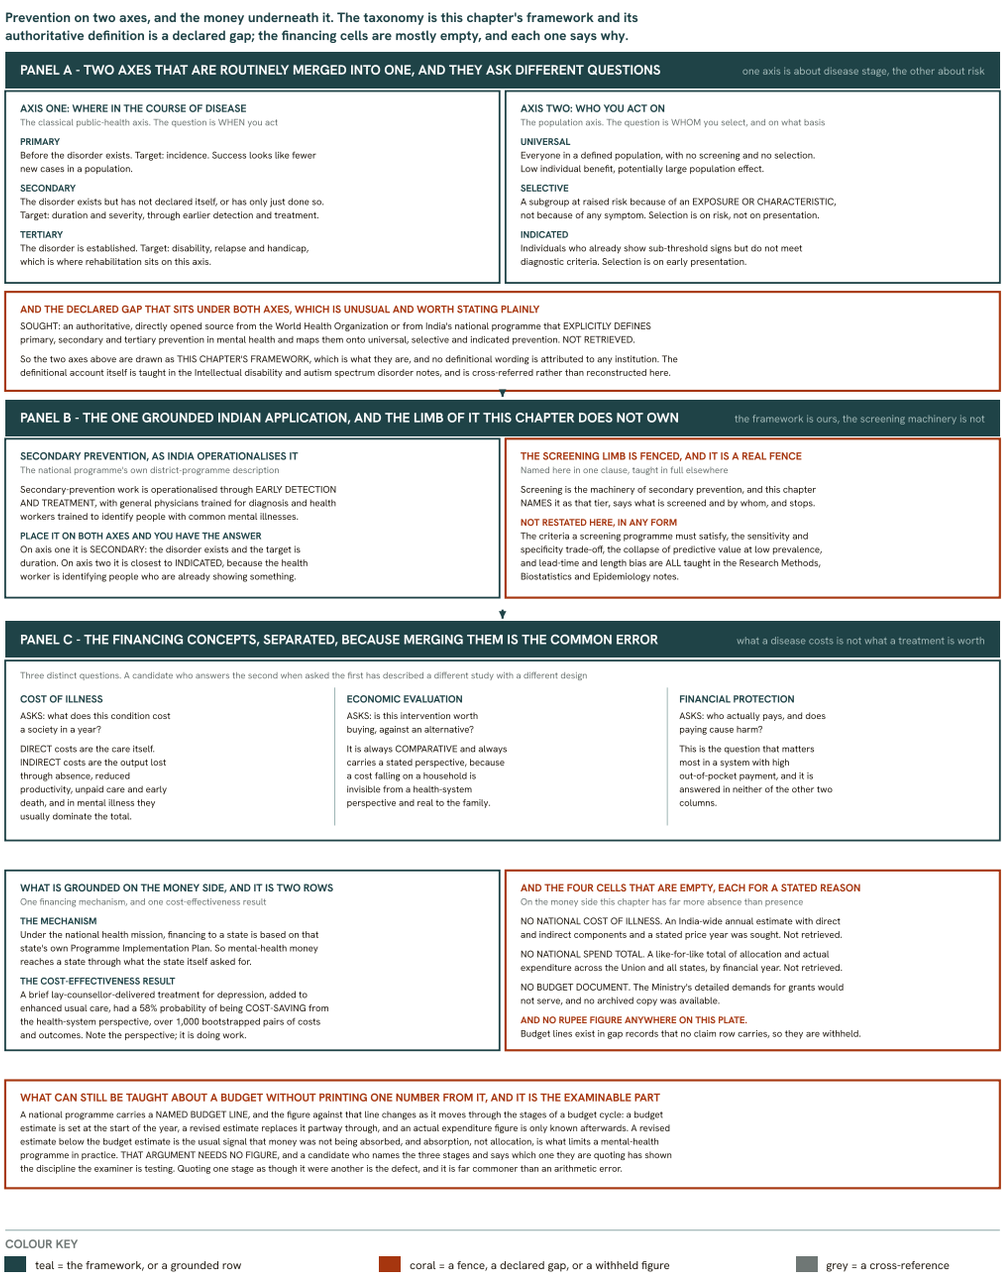


Fig 39.12 Prevention on two axes, and the financing underneath, with the empty cells labelled. The two axes ask different questions, disease stage against risk basis, and they are drawn as this chapter's framework because an authoritative source explicitly defining primary, secondary and tertiary prevention in mental health and mapping them onto universal, selective and indicated prevention was sought and not retrieved. The one grounded Indian application is secondary prevention operationalised as early detection and treatment through trained general physicians and identifying health workers; the screening machinery under it is fenced to the Research Methods, Biostatistics and Epidemiology notes. On financing, state money flows through the state's own Programme Implementation Plan, and a brief lay-counsellor-delivered treatment had a 58 per cent probability of being cost-saving from the health-system perspective across 1,000 bootstrapped pairs. No rupee figure appears: no national cost-of-illness estimate and no like-for-like national spend total were retrieved, and the budget lines that exist sit in gap records no claim row carries, so the budget argument is taught through its stages instead.

34.3 The recipient axis, and what it changes about delivery

The second axis is about selection, and its practical value is that it predicts the delivery problem before the intervention is designed. A universal intervention has no selection step at all, so it needs no case-finding apparatus and no consent to being identified as at risk, but it must be cheap per head because everybody receives it and it must be acceptable to people who will not benefit. A selective intervention requires a defensible basis for defining the subgroup, which means either a routinely available marker or a screening step, and it carries a labelling cost that the universal version does not. An indicated intervention requires individual identification, which is the most expensive selection of all and generates the largest ethical obligations, because a person identified as showing early signs is now known to the service and must be offered something.

The examinable consequence is that the cost profile inverts along this axis. Universal interventions have the lowest unit cost and the highest total cost. Indicated interventions have the highest unit cost and reach the fewest people, but concentrate the benefit where risk is highest. Neither is superior in principle, and a programme that has chosen one without stating why has usually chosen by convenience.

Whole population	**Raised-risk subgroup**	**Early signs, identified**	**Established disorder**
WHO A UNIVERSAL INTERVENTION IS OFFERED TO	WHO A SELECTIVE INTERVENTION IS OFFERED TO	WHO AN INDICATED INTERVENTION IS OFFERED TO	WHO SITS OUTSIDE PREVENTION ALTOGETHER

The fourth tile is there deliberately. People with an established disorder are not the recipients of prevention on the second axis at all; they are the recipients of treatment, and of tertiary prevention on the first axis. Keeping that tile in view is what stops a service describing its ordinary clinical work as prevention because the word attracts funding.

34.4 The routing table, which is what this section is for

Every substantial preventive activity in community psychiatry is taught somewhere, and the taxonomy's job is to say where. The table below is the chapter's routing device: it places each application on the axes and names where it is developed, so that the argument is made once, in full, in the right place.

APPLICATION	TEMPORAL TIER	RECIPIENT AXIS	WHERE IT IS DEVELOPED
Restricting access to lethal means	Primary	Universal	The means-restriction and gatekeeper material of these notes, with the population argument in the Somatic-symptom, sexual, impulse-control disorders and suicide notes
Responsible media practice on suicide	Primary	Universal	The means-restriction and gatekeeper material of these notes

APPLICATION	TEMPORAL TIER	RECIPIENT AXIS	WHERE IT IS DEVELOPED
Gatekeeper training for non-clinicians	Secondary	Selective, by role rather than by risk	The means-restriction and gatekeeper material of these notes
A national suicide prevention strategy as a whole	Spans all three	Spans all three	The national suicide prevention material of these notes; the three-tier grid worked through suicide exemplars is in the Somatic-symptom, sexual, impulse-control disorders and suicide notes
Anti-stigma work aimed at the general public	Primary, acting on a determinant of access	Universal	The anti-stigma intervention material of these notes
Mental health literacy and awareness campaigns	Primary and secondary, since recognition shortens delay	Universal	The mental health literacy material of these notes
School-based programmes	Primary, with an indicated limb inside them	All three, in one setting	The school and adolescent material of these notes; the Indian platforms are in the Child psychiatry: assessment, treatment, services and protection notes
Workplace programmes	Primary through tertiary, depending on the component	All three	The workplace material of these notes
Early detection by trained general doctors	Secondary	Indicated, on presentation	This section, and the primary-care integration material of these notes
Screening pathway design in one population	Secondary	Universal or selective, by design	The Perinatal/women's mental health and emergency psychiatry notes
The methodology of screening and its biases	Applies to secondary	Not applicable	The Research Methods, Biostatistics and Epidemiology notes
Rehabilitation, supported employment and community services	Tertiary	Established disorder, so outside the recipient axis	The rehabilitation and community-care material of these notes

RULE **Placing an intervention on the two axes tells you what evidence it needs and what denominator it must report against; it tells you nothing about whether it works.** A tier is a description, not a recommendation. Every row above still has to be argued on its own evidence in the place where it is developed, and the taxonomy's only claim is that the argument belongs in one place.

34.5 How the Indian programme operationalises the secondary tier

The abstraction becomes concrete at exactly one point in the Indian system, and it is worth stating precisely because it is the answer to the most likely viva question on this topic. In **India's National Mental Health Programme**, the secondary tier is operationalised as detecting illness early and treating it, with general physicians trained to diagnose and treat and health workers trained to identify.

Read the structure of that arrangement, because it explains the whole design. The Indian secondary tier is delivered not by an instrument but by a workforce: the detection step is a trained general doctor or health worker recognising a presentation in the course of ordinary contact, and the treatment step follows immediately from the same training. There is no separate screening programme with an invited population and a defined interval; there is an opportunistic detection system riding on existing clinical contact. That is a defensible design in a system with limited specialist capacity, and it has a specific consequence: the capacity of the secondary tier is set by training throughput and by supervision, not by procurement.

IN INDIA

Because the national programme puts its secondary tier in the hands of trained general physicians and identifying health workers rather than in a screening apparatus, a district that wants to strengthen this tier should be counting how many of its general doctors and frontline workers have actually been trained, how recently, and whether anyone supervises them afterwards. Buying or adopting an identification tool changes nothing if the person holding it has no training and no route for a positive finding. Training throughput and supervision are the levers this design exposes, and they are the ones a district can actually pull.

PITFALL

Treating the training as complete once it has been delivered. Detection skill in non-specialists decays without reinforcement, and the trained doctor is transferred, promoted or retires while the programme's records continue to count them. A district reporting cumulative numbers trained is reporting a historical total, not a present capacity, and the gap between those two quantities widens every year the programme runs.

34.6 Tertiary prevention, and why it loses the argument it should win

The largest preventable burden in psychiatry is disability in people who already have a diagnosis, and almost nobody calls the work that reduces it prevention. Relapse prevention, continuity of medication, family support, supported housing, supported employment and the

maintenance of social role are all tertiary prevention, and each prevents a specific, countable harm.

- A readmission that did not happen, which is the outcome most services already record and least often claim.
- A job lost during a relapse, and with it the income and the routine that supported recovery.
- A tenancy or a place in the family home lost while somebody was too unwell to hold it.
- A relationship that ended because nobody supported the family through an episode.
- A contact with the criminal justice system that followed from untreated illness rather than from anything else.

Each item on that list is countable, is already recorded somewhere, and has a cost attached to it in some department's accounts.

The reason this limb is systematically under-argued is partly linguistic and partly budgetary. Linguistically, prevention is heard as meaning stopping something from starting, so activity in people who are already ill is coded as treatment and competes for treatment budgets rather than for prevention ones. Budgetarily, the benefits appear as costs that did not occur, in a budget line different from the one that paid for the service, and often in a different department altogether. A supported housing placement that prevents several admissions saves money in the hospital budget and spends it in the community one, so the service that produced the saving does not receive it.

The counter-argument a candidate should be able to make is that the tertiary tier is the only one whose effect can be demonstrated in a defined denominator over a short period, because the population is already enumerated and the outcomes are already recorded. Where a primary-prevention claim requires a large population and a long follow-up, a tertiary-prevention claim can be tested on a service's own caseload against its own readmission record. The economics of that argument, and the methods for making it, are developed in the mental health economics and financing material of these notes.

34.7 What this chapter cannot establish, and how to answer without it

One thing an authoritative section on this topic would ordinarily supply is declared here as a gap. No authoritative source could be opened that explicitly defines primary, secondary and tertiary prevention in mental health and maps them onto the universal, selective and indicated scheme in a single statement of the taxonomy. What was sought was exactly that: an international agency publication on the prevention of mental disorders, several of its repository records, secondary reproductions of the taxonomy in the indexed literature, and Indian operational guidance for mental health at the primary-care level. The agency records returned access errors, the literature routes returned verification pages rather than article text, and the Indian operational documents supplied promotion, detection, screening and referral examples without labelling the full taxonomy.

The honest consequence is stated rather than concealed: the taxonomy is taught here as the standard framework it is, its definitions are cross-referred to the chapter that carries them, and no issuing body or year is attached to it in this section. A candidate asked to attribute the framework should say which chapter of their own reading carries the definitions and should not attach an agency and a date from memory, because a framework misattributed to a body that did not issue it is a discrediting error in an otherwise correct answer.

TRAP **Treating the tier label as though it were evidence of effect.** Candidates write that an activity is primary prevention and then proceed as though its benefit had been established by the classification. It has not: the tier states when and on whom the intervention acts, and every claim about whether it works is a separate claim requiring its own evidence, its own population and its own outcome. The most common form of this error is asserting that a public awareness activity prevents disorder, when the evidence available concerns knowledge or attitudes and the outcome asserted is incidence.

The two axes answer different questions, when and to whom, and the taxonomy's job in a community psychiatry chapter is to route each application to the place where its evidence is argued in full.

35 · Mental health economics: cost of illness, cost-effectiveness and financing

The argument about what mental health should cost is made whether or not psychiatrists take part in it, and it is settled by people who read budget documents rather than case notes. That is the practical reason this material is examinable: a clinician who cannot say what a service costs, what it saves, who pays and where the money enters the system is unable to argue for it in the only language the argument is conducted in. The topic divides cleanly into a small number of questions that are frequently run together and should not be. What does this illness cost, and to whom. Is this intervention worth its cost compared with the alternative. How does money actually reach a mental-health service. And what happens between an allocation and an expenditure. The classic economic argument for funding a general-hospital liaison service, with its historical figures and the observation that Indian cost-effectiveness evidence for such services is what is missing, is already made in the Consultation-liaison, geriatric and transcultural psychiatry notes and is not repeated or re-declared here. Insurance benefit design, parity and the regulator's mandate belong to the insurance and financial-protection material of these notes.

35.1 Cost of illness, and what the total is and is not

A **cost of illness** study answers a descriptive question: what does this condition consume, across a defined population, over a defined period. It is a burden statement, not an evaluation, and it makes no comparison between courses of action. Its components separate as follows, and the separation matters because different components fall on different payers.

COMPONENT	WHAT IT COUNTS	WHO USUALLY BEARS IT
Direct medical	Consultations, medicines, investigations, admissions, procedures	The health system, the household, or an insurer, in a proportion that varies by country
Direct non-medical	Transport to care, food and lodging for an accompanying relative, informal payments	The household, almost entirely, and it is routinely omitted from health-system costings
Indirect	Lost work and productivity of the patient, and of family members providing care	The household and the wider economy; the largest component for most mental disorders
Intangible	Pain, distress and loss of quality of life not otherwise captured	The patient and family; usually acknowledged and not costed

Two features of a cost of illness total decide whether it can be used at all. The first is the **perspective**: a total computed from the health system's point of view counts what the health system spends, and a total computed from a societal perspective adds the household's costs and the lost production. These give very different numbers for the same disease, and in mental illness the difference is unusually large because the indirect component dominates. A figure quoted without its perspective is uninterpretable.

The second is the counterfactual, and this is where the figure is most often misused. A cost of illness total is not a sum that could be saved. It is what the condition currently consumes, and no intervention removes all of it; some of it would be consumed by whatever the person had instead, and some of it is the cost of the treatment that keeps them well. The only defensible use of the total is to establish scale and to argue that a problem consuming this much deserves attention, and the only way to establish a saving is a comparison between courses of action, which is a different study.

TRAP **Presenting a cost of illness figure as the money a proposed programme would save.** The two answer different questions and the first cannot substitute for the second. A burden total has no comparator, so it cannot support any claim about a difference; and treating the whole burden as recoverable implies an intervention that abolishes the disorder at no cost, which nothing does. If the question is whether a programme is worth funding, the answer requires an economic evaluation and not a burden estimate.

35.2 What India's cost of illness figure is, and why it is not printed here

An examiner may reasonably expect an Indian national figure at this point, and the honest answer is that this chapter does not have one. What was sought was an official India-wide annual cost of illness estimate for mental disorders, with its direct and indirect components separated and a stated price or reference year. The national survey's summary report was consulted, the detailed report's official address returned an error, the ministry's mental-health and programme pages were searched, and indexed-literature and international-agency routes were tried for an Indian economic estimate. No official national estimate meeting that description was obtained.

The absence is worth teaching rather than concealing, and it is not a small one. Without an official national burden estimate, an Indian argument for mental-health investment has to be assembled from international burden estimates, from disability measures whose derivation belongs to the burden material of these notes, and from individual studies with local populations. Each of those is defensible and none of them is what a finance ministry asks for, which is one figure attributable to a national body with a stated year. A candidate who names the construct, says the national estimate is not established, and explains what the argument has to be built from instead has given a better answer than one who quotes a figure whose issuing body they cannot name.

35.3 Economic evaluation, and the four forms it takes

An **economic evaluation** always compares at least two courses of action on both their costs and their consequences. If it compares costs only, it is a costing exercise. If it compares consequences only, it is a clinical study. The forms differ in one respect only, which is how the consequence is valued, and that difference determines what question each can answer.

FORM	HOW THE OUTCOME IS VALUED	THE QUESTION IT ANSWERS	ITS LIMITATION
Cost-minimisation	Not valued; outcomes are assumed equivalent	Which of two equally effective options is cheaper	The equivalence assumption is rarely justified and is often asserted rather than shown
Cost-effectiveness	In natural clinical units, such as symptom-free time or cases detected	What it costs to buy one more unit of this particular outcome	Cannot compare across conditions, because the units differ
Cost-utility	In a combined measure of length and quality of life	What it costs to buy a comparable unit of health, whatever the condition	Depends on how the quality weights were derived and in which population
Cost-benefit	In money, including the health outcome itself	Whether total benefits exceed total costs, in one currency	Requires putting a money value on health, which is contested and rarely done well

The practical significance for mental health is that the middle two dominate, and the choice between them is a choice about who the result must persuade. A **cost-effectiveness** result in natural units speaks to a psychiatrist and to a programme manager who already accept that the condition matters. A cost-utility result is the one that lets mental health compete for a share of a health budget against cardiology and oncology, because it converts the outcome into a currency the whole budget uses. A specialty that only ever produces results in its own units is arguing where nobody else is listening.

35.4 Reading a cost-effectiveness result without being deceived by it

A cost-effectiveness figure carries a set of design choices, each of which can move it substantially, and each of which is a legitimate examination question.

- The perspective, again: a result from the health system's point of view ignores what the household paid to attend, and interventions that shift cost onto families look efficient from inside the system.
- The comparator: an intervention compared with nothing will almost always look good, and the informative comparison is with what the service would otherwise have done.
- The time horizon: benefits that accrue late are missed by a short study, and mental-health interventions with relapse-prevention effects are systematically undervalued by short follow-up.
- Discounting: costs and benefits arriving at different times are not equivalent, and the rate chosen affects any intervention whose costs are immediate and whose benefits are later, which describes most of prevention.
- Uncertainty: a single point estimate hides the range of plausible results, so a defensible evaluation reports how likely the conclusion is to hold rather than asserting it.

An Indian result illustrates the last of these particularly well. In a trial-based economic evaluation conducted in primary health centres in Goa, the programme under evaluation, identified in the trial by the acronym **HAP**, delivered alongside enhanced usual care, had a **58 per cent** chance of being cost-saving from a health system perspective.

Notice the form of that statement, because the form is the teaching. It is not a claim that the programme saved money; it is a statement of how often the intervention came out cheaper across repeated resampling of the paired cost and outcome data. A probability just above the halfway mark is a genuinely useful result and an honest one: it says the intervention is about as likely as not to pay for itself in health-system costs, while delivering its clinical benefit. Reported as a point claim that the programme is cost-saving, the same analysis would be a misrepresentation, and reported as a failure to demonstrate saving it would be a different one.

58 per cent	Not established	Not retrieved	Withheld
CHANCE OF BEING COST-SAVING, HEALTH-SYSTEM PERSPECTIVE, ONE GOA PRIMARY-CARE TRIAL	OFFICIAL INDIA-WIDE ANNUAL COST OF ILLNESS FOR MENTAL DISORDERS	LIKE-FOR-LIKE NATIONAL ALLOCATION AGAINST ACTUAL EXPENDITURE	SCHEME-LEVEL BUDGET LINES NO CLAIM ROW HERE ASSERTS

RULE **A cost or cost-effectiveness figure means nothing without its perspective, its comparator and its time horizon, and a probability of being cost-saving is not a statement that it saved money.** Those four attributes are the whole of the appraisal, and a figure quoted without them is the economic equivalent of a prevalence quoted without its population.

35.5 How money actually reaches a mental-health service

Financing mechanisms are the least glamorous part of this topic and the part that most determines what a district can do. The generic mechanisms are few. General taxation funds a service directly from government revenue. Social or mandatory health insurance funds it from contributions. Voluntary or private insurance funds it from premiums. Out-of-pocket payment funds it from the household at the point of use. External or donor financing funds it from outside

the country. Most systems, India's included, use several at once, and the mix determines who is protected and who is not.

What matters clinically is that these mechanisms differ in when the money is paid relative to when the illness occurs. Pre-paid mechanisms, whether tax or insurance, spread the cost across a population and across time. Out-of-pocket payment concentrates it on the household at the moment it is least able to bear it, which for a chronic mental illness means repeatedly, over years, in a family whose earning capacity the illness has already reduced. That is the whole argument for financial protection, and it is an argument about the timing of payment rather than about its total.

The Indian route by which programme money reaches a state has a specific and consequential shape. Under the **National Health Mission**, money reaches a state against that state's own Programme Implementation Plan. The plan is the state's own proposal of what it intends to do, and the money follows the proposal.

IN INDIA

Because national mission money reaches a state against its own Programme Implementation Plan, the operative lever for anyone trying to fund a mental-health activity in a district is what gets written into that plan before the financial year, not what is argued for afterwards. An activity nobody proposed is an activity nobody funded, and a clinician who wants a post, a training programme or a drug budget has a specific and dated point of entry: the plan-writing cycle. This is the single most actionable piece of financing knowledge in the topic, and it is administrative rather than economic.

35.6 Allocation against expenditure, and the gap between them

A budget line is a permission to spend, not a record of spending, and in mental health the two diverge more than in most programmes. The reasons are structural and are worth knowing because they are what an examiner is testing when they ask why increased allocation has not produced increased service. Money allocated to posts cannot be spent if the posts cannot be recruited to, and specialist mental-health posts are among the hardest to fill. Money allocated late in a financial year cannot be spent before it lapses. Money allocated for capital cannot be moved to recurring costs even when recurring costs are the constraint. And money released against a plan that was written optimistically returns unspent, which then becomes an argument for allocating less next time.

That last mechanism is the important one and it is self-reinforcing. Under-spending is read by a finance authority as evidence of low absorptive capacity, which justifies a smaller allocation, which reduces the capacity further. Breaking that loop requires spending on the things that can actually be delivered, which usually means training, supervision and recurring supply rather than infrastructure, and it requires the plan to be written by somebody who knows what the district can absorb.

PITFALL

Quoting an allocation figure as though it described a service. An allocation is what a document says may be spent; it is not what was released, not what was spent, and not what reached a patient. The chain from allocation to release to expenditure to utilisation loses value at each join, and a claim about service delivery supported only by a budget number has skipped every one of them.

35.7 The financing figures this chapter deliberately does not print

This section is where a chapter is most tempted to supply remembered numbers, and several that would ordinarily appear are declared here as gaps instead. Each is a real construct with a real answer somewhere; none is established by anything this chapter could open and verify, and a figure written from recollection into a financing argument is the kind of error that survives into somebody's answer script.

No like-for-like national total of mental-health allocation and actual expenditure, across the union and all states and union territories with each financial year identified, was obtained. Union budget documents, expenditure documents, the ministry's annual report, the national programme's own page, state plan and record-of-proceedings indexes, and the programme's framework and plan guidelines were all approached. The pieces exist in different documents on different bases and were not reconcilable into one comparable total.

The national programme's own line in the detailed demands for grants could not be reached: the address that would carry it returns an error after a domain migration, a repair attempt returned the same dead address, and no archived snapshot was available. The scheme-level budget lines for the national tele-mental-health programme, across recent actual, budget and revised estimates, and the plan-period allocations to the district programme, are likewise not printed here: those figures exist in budget documents but no claim row in this chapter's ledger asserts them, so they are withheld deliberately rather than reconstructed.

The teaching that survives all of that is the discipline itself, and it is the discipline the whole chapter is built on. An Indian financing figure is a fact only when it carries the document that issued it, the financial year it refers to, whether it is an allocation, a revised estimate or an actual expenditure, and what it is a proportion of. Those four attributes are exactly what makes budget figures so easy to misquote: the same programme in the same year has several different correct numbers attached to it, and quoting the wrong one confidently is indistinguishable, to a reader, from quoting an invented one.

35.8 Financial protection, and what a household actually pays

The last question is what the illness costs the family, and it is the one a clinician meets most directly. Out-of-pocket payment for a chronic mental illness is not a single event but a stream: repeated medicine purchases, repeated travel, repeated days of lost earning for the patient and for whoever accompanies them, and, where the public service is unavailable or unacceptable, private consultation fees. The stream continues during the periods when the person is well, which is exactly when a family is most likely to stop paying it, and non-adherence for cost reasons is often invisible to a clinician who never asks about the transport fare.

The national survey did produce a household expenditure finding for care of a person with a mental disorder, together with an observation about the absence of state or insurance coverage for most families. That finding is a declared gap in this chapter and its figure is not printed, because no claim row here asserts it. The construct is named without its rupee amount, and the reasoning around it stands on its own: an expenditure of this kind, incurred monthly and indefinitely by a household whose earning capacity the illness has reduced, is a financial-protection failure whatever its exact size.

The specific Indian mechanisms of financial protection, meaning the public insurance scheme, the parity requirement for mental illness and the regulator's mandate, are taught in the insurance and financial-protection material of these notes and are not restated here. What belongs to this section is the economic point behind them: financial protection is achieved by moving payment from the point of use to a pre-paid pool, and any scheme is judged on whether the household's payment at the point of use actually falls, not on whether a benefit exists on paper.

> A cost of illness total states scale and can never state a saving; only an economic evaluation with a named perspective, comparator and time horizon can do that, and in India the money reaches a service through the state's own plan.

36 · Digital mental health interventions and apps in the Indian context

The interesting question about a mental-health app in India is almost never whether it works, and almost always whether anybody who needs it can reach it, will keep using it, and has somebody at the other end who acts on what it produces. That is the difference between the product question and the system question, and it is the boundary this section works to. Whether a given digital therapeutic has established efficacy, what trial standards it must meet, how software is classified for regulatory purposes and what it means for such a product to be prescribable are taught in the Recent advances and landmark studies notes. What belongs here is the Indian implementation picture: access and the divide inside it, adoption and adherence, integration with an existing service, equity, governance and programme monitoring. Remote consultation between a patient and a clinician is a service model in its own right and is taught in the telepsychiatry material of these notes; this section covers what the health system does with software rather than what it does with a video link.

36.1 A functional taxonomy, because the technology label explains nothing

Grouping digital tools by their platform, meaning apps against websites against messaging, tells you nothing useful. Grouping them by what function they perform in a service tells you who has to act, what has to exist behind them, and how they fail.

FUNCTION	WHAT IT DOES	WHO MUST ACT FOR IT TO MATTER	HOW IT CHARACTERISTICALLY FAILS
Information and psychoeducation	Explains a condition, a treatment or a service to a public or a patient	Nobody, in the short term; it is a one-way transfer	Written in a register and a language its intended audience does not use
Self-guided intervention	Delivers a structured therapeutic programme with no clinician involved	The user alone, repeatedly, without prompting	Attrition, which is very high without human contact
Guided or blended intervention	Delivers the same programme with a person supporting adherence	A supporter, who need not be a specialist	The support component is dropped first when the budget tightens
Clinician-facing decision and workflow support	Prompts assessment, records outcomes, flags who is not improving	The clinician, and whoever governs the workflow	Added to a workload with no time released, so it is bypassed
Screening and case-finding by a frontline worker	Structures an enquiry the worker is already making	The worker, and the service receiving the positives	Positives are generated with no destination to send them to
Monitoring and follow-up	Collects symptoms, adherence or attendance between contacts	Somebody who reads it and responds	Data accumulate that nobody has been made responsible for

The fourth column is the section in miniature. Every failure listed is organisational rather than technical, and none of them is fixed by a better application. That is why an Indian digital mental-health question is a service question.

RULE **A digital tool is worth what the response behind it is worth.** Software that identifies distress, records a symptom score or flags deterioration changes nothing unless a named person is obliged to act on the output within a defined period. Where no such person exists, the tool has converted a clinical problem into a data problem.

36.2 Access, and the divide inside the population that most needs the service

The **digital divide** is the reason this matters. Digital delivery is routinely presented as an access solution, and it can be, but only against an explicit account of who cannot use it. The entry requirements are specific and each is distributed unevenly in the same direction as the underlying disadvantage.

- A device that the person controls, rather than a household device controlled by somebody else.
- Connectivity that is reliable enough and cheap enough for repeated rather than occasional use.
- The language of the interface, which for a person in distress must be the language they think in and not merely one they can read.

- Enough literacy, and enough familiarity with the conventions of an interface, to navigate it without help.
- Privacy at the moment of use, which is the requirement most often forgotten and the one hardest to meet in a crowded household.

The shared-device problem deserves separate emphasis because it inverts the usual assumption. Where a phone is shared within a household, an application on it is visible to whoever else uses it, and a browsing history, a notification or an installed icon can disclose a person's help-seeking to the very people they may need to keep it from. For a woman in a household where help-seeking is contested, and for an adolescent, a digital route can be less private than a clinic visit rather than more. A service that assumes digital equals confidential has assumed the opposite of what may be true.

> **PITFALL**
>
> Building in one language and adding others later. Language is treated as a translation task to be done after the product works, so the pilot runs in English or in the state's dominant language, the evaluation reports on the population that could use it, and the populations excluded never appear in the denominator. Because the excluded groups are also the underserved ones, the programme's own evaluation will report success while the access gap it was built to close is unchanged.

36.3 The Indian evidence, and what it points to

The strongest Indian signal for digital mental health does not come from a consumer product but from a multifaceted intervention delivered inside a rural primary-care system. In rural primary-health-centre clusters in Haryana and Andhra Pradesh, among adults at high risk of depression or self-harm, those who received the intervention had lower depression scores than those receiving usual care, and nearly every secondary outcome moved in the same direction. The outcomes examined alongside depression included anxiety, remission, access to treatment, knowledge, attitudes, stigma and intended help-seeking.

Two features of that result are worth extracting, because they are what an Indian answer should be built on. The first is that the digital component sat inside a service rather than beside it: the technology supported frontline health workers and primary-care clinicians who were already in contact with the population, which means the entry requirements listed above were largely met by the service and not by the patient. The second is the breadth of the outcome set. An intervention that moves knowledge, attitudes, stigma and intended help-seeking as well as symptoms is acting on the pathway to care and not only on the illness, which is the correct ambition for a population intervention and is not something a standalone self-guided product is designed to do.

IN INDIA

Because the Indian evidence that has actually moved depression scores comes from an intervention delivered through rural primary-care clusters, associated with lower depression scores than usual care, the practical decision for a district is to treat a digital tool as a component bolted into an existing pathway with named workers behind it, not as a service that can be procured and pointed at a population. The corollary is a procurement rule: before adopting any tool, name the worker who will use it, the time released for them to do so, and the clinician who receives what it produces. A tool adopted without those three has been bought rather than implemented.

36.4 Adoption and adherence, which are separate problems

Digital interventions have a characteristic and well-described failure curve, and it has a shape that no other intervention type shares. **Uptake** is easy and cheap, because installing something costs the user almost nothing. **Continued use** collapses quickly, because stopping also costs almost nothing. The result is that the population who complete a self-guided programme is a small and highly selected subset of the population who started it, and it is selected for exactly the characteristics that predict good outcome anyway.

The single most reliable modifier of that curve is human contact. A brief, scheduled contact from a person, who need not be a mental-health specialist and need not deliver any therapy, substantially changes the proportion who continue. That is the entire logic of guided digital intervention, and it also explains why the guided version costs more than the self-guided one and is worth it: the support is not an accessory to the software, it is the component that makes the software reach anybody.

MNEMONIC

Access, Adoption, Adherence, Action

The four gates a digital intervention has to pass, each with its own denominator and its own failure. Who could reach it. Who started it. Who kept going. And who did something with what it produced. A programme reporting only the second gate has reported the easiest number it has.

Who could reach it	Who started	Who continued	Who was acted on
ACCESS, DENOMINATOR IS THE INTENDED POPULATION	ADOPTION, DENOMINATOR IS THOSE WHO COULD REACH IT	ADHERENCE, DENOMINATOR IS THOSE WHO STARTED	ACTION, DENOMINATOR IS THOSE WHO PRODUCED A SIGNAL

36.5 Integration, or why an application is not a service

An application becomes part of a health service at a small number of specific joins, and naming them is the most useful practical content in this topic. Who introduces the tool to the patient, and at which contact. Whether what the tool records enters the clinical record or sits in a separate system nobody opens. Who is responsible for a concerning response, within what period, and by what route. What happens when the person stops using it, since disengagement from a monitoring tool is itself clinical information. And who supports the user when the tool does not work, which in practice is whoever they saw last.

Safety escalation is the join with the highest stakes and the one most often unspecified. A tool that asks about self-harm will receive positive answers, at any hour, from people whose location and identity it may not know. A programme that has not decided in advance what happens then has built a mechanism for collecting disclosures it cannot respond to. The decision is not primarily technical: it is about who is on duty, what they are permitted to do, and what the local escalation route is.

36.6 Governance, data and the questions an Indian programme has to answer for itself

Digital mental-health data are among the most sensitive categories of personal information, and the governance questions are concrete rather than abstract. Where are the data stored and under whose control. Who may access them, and can an employer, an insurer or a family member obtain them. What was the person actually told, and was consent obtained in a language and a form they understood rather than by scrolling past a screen. What happens to the data if the provider is acquired, changes its terms, or closes. And who is accountable when the tool gives advice that turns out to be wrong.

This chapter carries no Indian instrument that settles those questions for mental-health applications, and none is attributed here. That absence should be stated as such rather than filled: a candidate who names the questions, says that a specific national framework for evaluating, endorsing or governing mental-health applications is not something they can attribute, and points to the chapter that carries the regulatory classification of software has answered better than one who names a framework from recollection. The regulatory and product-classification material sits in the Recent advances and landmark studies notes.

One governance point can be made without any instrument, because it follows from the service logic already established. A public programme that recommends a third-party application to a population has taken on responsibility for it, and has done so without any of the assurance machinery it would apply to a medicine or a device. That is a decision worth making explicitly rather than by default, and the minimum discipline is to know who evaluated the tool, against what, and who reviews that judgement when the product changes.

36.7 Monitoring a digital programme, and the measures that flatter it

Digital programmes generate more data about themselves than any other kind of service, and almost all of it measures the wrong thing. Downloads, registrations, page views, sessions and average session duration are all easy to collect, all report on activity, and none has the intended population in its denominator. A programme reporting them is reporting the traffic its own promotion generated.

The measures that answer the question are the four gates named above, each reported against its own denominator, plus the clinical outcome in the population reached. Two further measures are worth insisting on because they are the ones that reveal an equity failure. First, the composition of the reached population compared with the intended one, broken down at least by sex, by rural or urban residence and by language, since a programme can look successful while serving a demographic that was never the target. Second, the proportion of signals generated that received

a response within the intended period, which is the measure that distinguishes a service from a data-collection exercise.

TRAP **Quoting downloads or registrations as evidence of a digital mental-health programme's reach.** A download is an installation event with no denominator, no population and no clinical content. It can be produced by advertising alone, it counts devices rather than people, and it says nothing about whether the application was opened twice. The reach question requires the intended population as a denominator, and the impact question requires an outcome measured in the people who actually used it.

36.8 What this section does not have, and how to answer without it

The figures a resident is most likely to be asked for here are the ones this chapter cannot supply, and the discipline of saying so is worth as much as the figures would be. No official national Indian figure for the reach, uptake or effectiveness of any mental-health application is carried here, and none is invented. No national count of Indian mental-health applications, of their users, or of their clinical outcomes is available to this chapter. No Indian figure for device ownership, connectivity or private access among people with mental illness is held either, which is why the access argument above is made structurally rather than quantified.

Nor is any prevalence figure supplied for online harms among Indian school-going children and adolescents. That figure is frequently quoted, the populations behind the various circulating versions differ, and a figure measured in one country's schoolchildren does not become an Indian figure by being applied to Indian ones. A service designing a response to online harm has to do so without knowing the size of the problem, which is uncomfortable and is the actual position.

There is a further consequence of that position which is worth drawing out, because it is the one a service actually has to act on. Where the size of a problem is unknown, a service cannot size its response to the problem, so it must size the response to what it can sustain and then measure what it meets. That is a different planning logic from the one used for a quantified problem, and it puts the emphasis on building a route that can absorb whatever arrives rather than on estimating demand in advance. Applied to online harm among young people it means an accessible, non-stigmatising point of contact and a workforce trained to receive the presentation, rather than a campaign scaled to a prevalence nobody has established.

What remains, and it is the durable part, is the framework: a digital intervention is characterised by its function rather than its platform, is limited first by access and then by adherence, is worth what the response behind it is worth, and is evaluated on the intended population as a denominator rather than on the traffic it generated. Those propositions hold whichever product is in front of you, and they are what an examiner is testing when they ask about digital mental health in the Indian context.

Access, adoption, adherence and action are separate gates with separate denominators, and the Indian evidence that has moved outcomes came from a tool embedded in a primary-care service rather than from a product handed to a population.

37 · Social determinants of mental health and global mental health movement

Almost every strong predictor of mental ill health in a population is something a psychiatrist cannot prescribe. Income, education, housing, safety at home, work, and the accumulated experience of childhood do more to shape the distribution of disorder in a district than the availability of any treatment does, and they are administered by departments the health service does not control. That is the uncomfortable structure of this topic, and it is what makes it a public-health subject rather than an aetiological one. The examinable content is not a list of hardships; it is the argument that follows from the list, which runs: the causes are largely social, the levers are largely outside health, therefore the health system's role includes evidence, advocacy and intersectoral work, and none of that reduces its obligation to treat the people in front of it. Migration and ethnicity taught as a clinical and epidemiological topic, with its mechanisms, its acculturation framework and its risk estimates, belongs to the Consultation-liaison, geriatric and transcultural psychiatry notes; migration appears here only as one determinant among others, without any of its figures.

37.1 What a social determinant is, and the life-course claim

A **social determinant** is a condition of life that influences health and is itself produced by how a society distributes resources, status and power. Two properties distinguish a determinant from an ordinary risk factor. It is structurally produced, so it is patterned across a population rather than randomly distributed. And it is modifiable by policy rather than by clinical action, which is precisely why it belongs in a chapter about programmes.

The organising proposition for the whole topic is that mental health, and much of the common mental disorder a population carries, is patterned by **the environments people live through**, meaning the social, the economic and the physical together, acting at different stages of a life rather than at one. Every clause of that proposition is doing work. It covers mental health and not only disorder, so the argument is about the distribution of wellbeing and not merely about caseness. It names three kinds of environment, so the determinants are not reducible to income. And the life-course clause is the decisive one: exposures act at different stages, they accumulate, and an exposure in early life can express itself decades later, which means the timing of an intervention matters as much as its content.

MNEMONIC

Cause, Course, Consequence, Contact

The four routes by which a social determinant reaches mental health, and they are separate claims requiring separate evidence. It can raise the chance that disorder occurs. It can worsen the course once it has. It can amplify the social damage the disorder does. And it can obstruct contact with a service, which is the access limb and belongs with the treatment gap material of these notes.

37.2 The determinants, grouped by the lever that would move them

Listing determinants is easy and unilluminating. Grouping them by which sector holds the lever converts the list into a plan, and it is the form an examiner rewards.

DETERMINANT	HOW IT ACTS ON MENTAL HEALTH	WHICH SECTOR HOLDS THE LEVER
Poverty and income insecurity	Chronic stress, debt, inability to sustain treatment, and constrained choice at every decision point	Finance, social protection and employment
Education	Shapes literacy, earning capacity and the ability to navigate services across the whole life course	Education
Adverse experience in childhood	Acts early, accumulates, and expresses itself long after the exposure has ended	Child protection, education, police and social welfare
Violence in the home and gender inequality	Direct exposure, and constrained ability to leave, to earn or to seek help independently	Home affairs, law, and women and child development
Housing and the physical environment	Crowding, insecurity of tenure, noise, heat, sanitation and the absence of private space	Housing, urban development and local government
Work and its conditions	Demand, control, security and reward, and the loss of all of them in unemployment	Labour, and employers
Alcohol and substance availability	Price and access shape consumption, and consumption shapes harm at population scale	Excise, revenue and law enforcement
Discrimination and social exclusion	Chronic exposure to devaluation, and reduced access to every other resource above	Law, education and every service that receives the public

The third column is the point of the table. Not one of these levers sits inside a mental-health budget, and most sit outside health altogether. A district mental-health plan that lists determinants without naming the department that owns each one has produced a description; a plan that names them has produced an addressee.



Fig 39.13 Determinants, culture and partnership: the layer upstream of every figure in this chapter, and the two limbs where the evidence runs out. The grounded determinant statement is that mental health and common mental disorders are shaped across the life course by social, economic and physical environments, in which the life-course clause, the three environments and the word shaped are each load-bearing. Global mental health began as a coverage argument: the call to action asked for scale-up of service coverage in all countries, especially low- and middle-income ones, and the Movement for Global Mental Health was launched. Culture is drawn only for its pathway consequence, culture-bound presentations approached through locally meaningful idioms of distress; the taxonomy, the named syndromes and the structured cultural assessment instrument are fenced to their owning chapters. Partnership carries one process finding, that trust, rapport-building and open dialogue are core requirements, while no primary source specifying bidirectional referral, healer training, safeguards, accountability and evaluation together was retrieved and the official Indian position is recorded as contradicting prior print rather than as merely absent.

37.3 The gradient, and what it implies for how a service is targeted

Determinants act as a **gradient** and not as a threshold. It is not the case that disadvantage below some line produces risk and everything above it is equivalent; the relationship runs across the whole range, so a service that concentrates entirely on the worst-off leaves most of the population-attributable burden untouched, because most people are not in the worst-off group.

That produces a design tension every programme in this chapter meets. Targeting the most disadvantaged is efficient per person reached and misses the bulk of cases. Universal provision reaches everyone and dilutes effort where need is lowest. The resolution used in practice is to provide universally while scaling the intensity of provision to need, so that everybody is covered and the most disadvantaged receive more. The principle carries a formal name in the health-inequalities literature, which this section does not attach because no claim row here carries it; the design instruction stands without the label, and a candidate should describe the arrangement rather than name a framework they cannot attribute.

RULE **A determinant is a cause the health system can usually only reach through another sector, and naming the sector is part of naming the determinant.** A statement that poverty causes mental ill health is not yet a public-health proposition. It becomes one when it says which policy instrument, held by which department, would move it, and what the health system's contribution to that argument is.

37.4 The nihilism trap, and the intervention menu that answers it

The commonest misuse of this material is to treat it as a counsel of despair: if the causes are social, clinical work is a distraction, and nothing short of redistribution will help. That inference is wrong in both directions and it is worth being able to dismantle it quickly.

It is wrong about causation, because a determinant raises the probability of disorder in a population without making the disorder in any individual less treatable. It is wrong about action, because the health system has several genuine roles that require no other department's permission. It can produce and publish the evidence that a determinant is acting, which is often the only quantitative case anybody has. It can advocate on the basis of that evidence in the forums where policy is set. It can design its own services so that they do not compound disadvantage, meaning free at the point of use, near where people live, in the language they speak, at hours that do not cost a day's wages. It can deliver interventions that buffer a determinant it cannot remove, such as support to families under strain. And it can treat, properly and without apology, the people the determinants have already made ill.

Set out as a sequence, those roles answer the objection completely, and they are worth being able to recite because they convert a general complaint into a work plan.

- Measure the determinant in the local population, since the health system frequently holds the only routine data on it.
- Publish and advocate, in the forum where the relevant policy is actually decided rather than in a clinical one.

- Remove the barriers the service itself creates, meaning cost, distance, language and appointment hours.
- Deliver the buffering interventions the health system can deliver alone, such as support to families under strain.
- Treat, at the standard the evidence supports, without letting social causation become a reason to expect less.

37.5 Intersectoral action, and what this chapter cannot attribute

Intersectoral action is the practical form of everything above, and it fails in predictable ways. It fails when it is convened without a shared measure, so each department reports its own activity and nobody can say whether anything changed. It fails when the health sector arrives with a request rather than with evidence and an offer. It fails when the benefits accrue to one department and the costs to another, which is the ordinary case. And it fails when it exists only at the level of a committee, with no counterpart at the district level where services actually meet.

An Indian national statement to this effect would be the natural anchor for this section, and it is declared here as a gap rather than supplied. The construct sought was the national survey's own systems framework proposition, that the determinants of mental health lie outside the health sector and that systems for care are incomplete if health-related sectors are not included. That proposition is not established here from a source this chapter opened and verified, so it is named as a construct and not asserted as a finding or attributed to a body. The reasoning above stands on its own as service reasoning; the national attribution does not.

IN INDIA

Because mental health and common mental disorders are shaped by social, economic and physical environments acting at different stages of life, a district mental-health plan that contains only clinical activities has, by construction, left its largest levers unaddressed. The practical decision is to require that the plan name, for each determinant it identifies, the department that holds the lever and the officer who will be approached, so that intersectoral intent becomes a set of addressees rather than a paragraph. A determinant with no named counterpart in another department will not be acted on, and writing it down is what converts the health system's evidence into somebody else's agenda item.

37.6 Migration, taken as one determinant among others

Internal migration belongs in this list because moving, whether across a border or within a country, rearranges almost every other determinant at once: income, housing, language, social network, entitlement and legal standing all change together, usually abruptly. That is why it behaves as a determinant rather than as a single exposure.

The boundary here is strict and is stated so a reader knows where to go. Migration and ethnicity taught as a clinical and epidemiological topic, with the competing mechanisms proposed to explain differences in rates, the framework used to describe adaptation after migration, and the estimates attached to any of it, are taught in the Consultation-liaison, geriatric and transcultural

psychiatry notes. This section prints no risk estimate for migrants in any generation, no Indian migration rate, and no estimate of any protective effect, and it attributes no position on migration to any Indian guideline.

What remains, and it is genuinely this section's, is the Indian internal picture as a service and equity problem. Internal migration within India moves large numbers of people between states and from rural to urban areas, and the service consequence is structural: entitlements, registrations and continuing prescriptions are commonly attached to a place of origin, while the person is now somewhere else. A worker who moves for seasonal or construction work loses continuity of any long-term treatment as a matter of administration rather than of clinical decision, and returns to their district when unwell. Any service designed for a settled population will systematically under-serve them.

The Indian evidence for that argument is declared as a gap. What was sought was an official Indian source documenting internal migration as a mental-health determinant together with its service-access or equity consequence, without reproducing the clinical and epidemiological migration material owned elsewhere. The national survey routes, ministry pages and international agency material approached returned general determinants or mentions of migration among proposed determinants, but no Indian estimate or official statement of the service consequence was retrieved. The structural argument above is therefore offered as reasoning about how entitlement systems work, and no Indian figure is attached to it.

PITFALL

Attaching a risk multiple to migration. The figures that circulate for psychotic illness in migrant populations come from particular host countries, particular migrant groups and particular ascertainment methods, and they are the subject of an active methodological argument taught in the chapter that owns the topic. Importing one into an Indian internal-migration argument compounds two errors at once: it transfers an estimate across populations that differ in almost every relevant respect, and it converts a service and equity argument into an aetiological claim that this chapter has no standing to make.

37.7 The global mental health movement, and what it actually argued

The **global mental health movement** is best understood as an argument rather than an organisation, and its central proposition is straightforward. A large share of the world's burden of mental disorder falls in countries with the fewest specialists, most people with a disorder in those countries receive no treatment, effective interventions exist and can be delivered by non-specialists, therefore the ethical and practical priority is to scale up coverage rather than to wait for specialist workforces to grow.

The movement's defining call, issued in **2007**, asked its readers to increase the coverage of services for mental disorders in every country and above all in low-income and middle-income ones. The wording contains the argument in miniature: what is demanded is *coverage*, which is a population measure with a denominator, and the demand is addressed to every country with a

particular emphasis rather than to poor countries alone. Shortly afterwards the Movement for Global Mental Health was launched as a network to carry that agenda.

2007	Coverage	Non-specialists	Not established here
CALL TO SCALE UP SERVICE COVERAGE, ALL COUNTRIES, ESPECIALLY LOW- AND MIDDLE-INCOME	THE MEASURE THE CALL DEMANDED, NOT SPECIALIST NUMBERS	THE WORKFORCE THE ARGUMENT RESTS ON	ANY INDIAN COVERAGE FIGURE FOR THIS ARGUMENT

The movement has standing objections, and a candidate who can state them fairly is demonstrating more than one who cannot. The first is that exporting a diagnostic and treatment framework developed in one part of the world risks displacing local understandings of distress and local ways of responding to it. The second is that scaling up treatment can absorb the entire agenda and leave the determinants untouched, which is the tension this whole section is about. The third is methodological: an intervention shown to work in one setting is not thereby shown to work in another, and the transportability of evidence is an empirical question rather than an assumption. The fourth is about who sets the agenda, since priorities, funding and publication have historically been concentrated in the countries with the fewest patients.

37.8 What the argument leaves an Indian service with

India occupies an unusual position in this debate, because it is simultaneously one of the settings the movement was arguing about and one of the places that produced the evidence the movement rests on. That gives an Indian clinician a stronger footing than the debate's usual framing allows: the question is not whether an external agenda should be accepted, but which parts of a largely home-grown evidence base a service should implement, and with what adaptation.

Three implications are worth carrying away. Coverage, not capacity, is the measure the argument is conducted in, so a service arguing for resources should be able to state what proportion of its intended population it reaches. Non-specialist delivery with specialist supervision is the mechanism the argument depends on, which is why it recurs in almost every section of this chapter. And the objections above are not reasons to abandon the programme; they are the design constraints it has to satisfy, meaning that a scaled-up service must be culturally intelligible, must be evaluated locally rather than assumed to transfer, and must not be allowed to become the whole of a country's mental-health policy while the determinants that fill it go unaddressed.

TRAP **Treating the social determinants as an alternative to services, or services as an alternative to the determinants.** Candidates pick one and argue against the other, and both arguments are weak. Determinants shape how many people become unwell; services determine what happens to those who do. Neither substitutes for the other, and the examinable position is that a health system owes evidence and advocacy on the first while delivering coverage on the second.

The determinants of mental health are largely produced outside the health sector and reached through other departments, and the global argument that follows is about coverage of a population rather than about the number of specialists a country has.

Culture, global mental health and traditional healing

38 · Cultural psychiatry, idioms of distress, culture-bound syndromes and explanatory models in the Indian context

Culture belongs in a public-health chapter because it decides where a population goes first, how long it stays there, and what it is holding when it finally arrives at a health service. That is a service fact with denominators and waiting times attached, and it is distinct from the cultural encounter at the bedside. The definitions of the cultural constructs, the umbrella that groups them and the classic roster of named syndromes are taught in the Psychotherapies, clinical assessment, MSE and nosology notes; the clinical encounter with a culture-named presentation, the negotiation between two accounts of one complaint and the placement of a cultural formulation inside a working assessment are taught in the Consultation-liaison, geriatric and transcultural psychiatry notes. This section takes none of that. It takes the population consequence: which door a community uses first, what the detour costs, what a general health service has to be able to receive, and what a service must build for any of this to change.

38.1 The pathway, read as a service problem rather than a curiosity

A **pathway to care** is the sequence of people a person consults between first noticing something is wrong and arriving at the service that will eventually treat them. Read at population scale it has three properties that matter to a planner and are invisible in the consulting room.

The first is that psychiatric services are almost never the first stop, and in many communities they are not the second or third either. The first consultation is usually inside the family, then with whoever the community treats as the appropriate authority for that kind of trouble, and then with a general practitioner or a general hospital. That is not a failure of the population; it is a rational sequence given how the problem was understood at the outset.

The second is that each step consumes time and money, and both are spent before any effective treatment starts. The cost is therefore paid twice over: once in the direct expenditure on consultations that did not help, and once in the untreated duration of the illness, which for several conditions is itself associated with worse outcome.

The third, and the one that turns this from description into planning, is that the pathway is modifiable. Where the population goes first is determined by what the community understands the problem to be, and by which doors are physically and socially available. Both can be worked on. A service that treats the pathway as fixed has accepted a delay it could have shortened.

RULE **The pathway to care is a property of the service system as much as of the culture, and it is measurable.** Where a population goes first is determined jointly by what it believes about the problem and by which doors are open, affordable and near. A service can change the second of those directly and the first only slowly, which is why pathway work begins with placing the service where the population already goes.

38.2 The detour, and what it actually costs a service

The **interval between onset and effective treatment** is the quantity that summarises all of this, and it is the number a service should want most and usually does not have. What can be said without any figure is what the delay is composed of, because each component has a different remedy.

STAGE OF THE PATHWAY	WHAT HOLDS THE PERSON THERE	THE LEVER A SERVICE ACTUALLY HAS
Inside the family, before any consultation	The problem is understood as something other than illness, or as something the family should manage	Population-facing explanation in the community's own vocabulary, which is the literacy and awareness work of these notes
With a community or faith-based helper	That helper is the recognised authority for this kind of trouble, is nearby, and is trusted	Partnership and referral arrangements, developed in the traditional and faith-based care material of these notes
With a general practitioner or general hospital	The complaint presented is one the general service can respond to on its own terms	Recognition capacity in general services, which is the primary-care integration work of these notes
Referred but not arrived	Distance, cost, stigma, or a referral with no named destination	Where the service is placed, what it costs, and whether the referral names a person and a date

Notice that only the first row is about belief. The remaining three are about proximity, trust, recognition and the mechanics of referral, all of which are ordinary service-design variables. It is a persistent error to treat the whole of the pathway delay as a cultural problem requiring cultural work, when most of its length is produced by where the service is and how hard it is to enter.

38.3 Idioms and explanatory models as population phenomena

A population does not present its distress in the categories a service uses to classify it. It presents in the vocabulary it already has, which may be bodily, social, moral or spiritual, and which is coherent within the community's own understanding of what has gone wrong. The taxonomy that groups these forms of expression, its definitions and the roster of named presentations that belong to it are taught in the Psychotherapies, clinical assessment, MSE and nosology notes, and the umbrella is named here once as the frame this argument rests on, without reproducing any part of it.

The stance this chapter takes toward those presentations is a service stance, and it has a source. In the Indian subcontinent, culture-bound presentations are more productively received as **locally meaningful idioms of distress** than as biomedical syndromes to be translated automatically at the door. Read as a service instruction rather than as a nosological position, that has concrete consequences.

- The presentation that arrives is data about the population's understanding, not noise to be filtered before the real assessment starts.
- A service that only recognises complaints already framed in its own categories will systematically under-detect, because the population is not obliged to speak in those categories.
- Population-facing communication that ignores the local vocabulary is speaking a language the audience does not use to think about the problem.
- An idiom carries information about severity and about what the person expects to happen next, which is why receiving it accurately affects what can be offered.

Two Indian presentations, dhat and koro, are the standard examples of a complaint whose form is shaped by a shared explanatory model, and they are named here only as that, examples of how a population's understanding routes a person to one door rather than another. Their phenomenology, their nosological status and their clinical handling are taught in the Somatic-symptom, sexual, impulse-control disorders and suicide notes and are not touched here.

IN INDIA

Because culture-bound presentations in this subcontinent are better received as locally meaningful idioms of distress than converted into a biomedical label at the door, an Indian outpatient service should build the recording of the person's own account of the problem into its first contact rather than leave it to whoever has time. The account is what the eventual explanation, the treatment plan and any psychoeducation have to connect to; a service that records only its own formulation has kept the half of the encounter it already understood and discarded the half that determines whether what it offers makes sense to the person receiving it.

38.4 Recognition in general health services, which is where most of it lands

The great majority of culturally framed distress presents to general health services and not to psychiatric ones, which makes **recognition capacity** in those services the single largest determinant of whether it is ever addressed. The design problem is specific: a general practitioner working under time pressure will act on the complaint as presented, and a complaint presented in bodily terms will generate an investigation rather than an enquiry.

The levers here are unremarkable and effective. Training aimed at recognising the common presenting forms in that population rather than at diagnostic subtlety. A short structured enquiry the clinician can actually complete in the time available. A referral destination that is named, near and reachable, since a referral into an abstraction is not a referral. And feedback to the referring clinician about what happened, which is the only mechanism that improves recognition over time and is almost universally absent.

PITFALL

Designing the service around the assumption that the health system is the first door. Outreach, awareness material and referral arrangements are all built as though the population's starting point were a health facility, so they address people who have already found the system. The population that has not is reached by none of them, and it is exactly the population whose pathway is longest. Finding out which door a community actually uses first, in that district, is a prior question and not a research luxury.

38.5 What a service must build for a cultural formulation to be usable at all

It is one thing to accept that a person's own account of their problem should inform their care, and another to run a service in which that happens reliably. The implementation questions are three, and they are the ones a programme has to answer.

The first is who is trained. If eliciting an account is treated as a specialist skill, it will happen only where specialists are, which is the smallest part of the system and the last stop on the pathway. If it is treated as a competence for whoever holds the first clinical contact, it happens where the population is, which requires a short, teachable version and a supervision structure behind it.

The second is at which contact it is applied. Applying it at every contact is unaffordable at scale; applying it at none makes the commitment decorative. The workable answer is to fix it to a defined point, usually the first assessment and any point at which the plan is not working, and to record the outcome where the next clinician will see it.

The third is what it costs in consultation time, and this is the question that decides whether the policy survives contact with a busy outpatient department. An enquiry that adds materially to every consultation in a clinic seeing very large numbers will be abandoned within weeks and will then be reported as having been implemented. The honest planning position is that time is the binding constraint, that the enquiry must be scoped to the time actually available, and that if the time does not exist the service should say so rather than mandate something it cannot deliver.

What such an enquiry is for is modest and specific. **An explanatory-model assessment** is aimed at three things a service otherwise never learns: what the problem means to the person, where they think help might come from, and what they expect a service to do. Each of those is directly actionable, which is why the enquiry is a service investment rather than an ethnographic one: the first tells you what the explanation must connect to, the second tells you who else is already involved, and the third tells you what will be experienced as help. The instrument built on this approach, its structure and its content are taught in the Psychotherapies, clinical assessment, MSE and nosology notes, and where the enquiry belongs inside a working clinical assessment is taught in the Consultation-liaison, geriatric and transcultural psychiatry notes.

MNEMONIC

Place, Partner, Practitioner, Proof

The four decisions that convert a cultural argument into a service change. Place the service where the population's pathway already runs. Partner with whoever holds the door it uses first. Train the practitioner who takes the first clinical contact, not only the specialist at the end. And prove it, by measuring where people came from and how long they took to arrive.

38.6 Measuring a pathway, which almost no service does

The whole of this argument becomes tractable the moment a service records where its patients came from. That is a small addition to a first-assessment record and it produces the only data that can tell a district where its access problem actually is.

First contact	Number of steps	Interval to treatment	Referral source
WHO THE PERSON CONSULTED BEFORE ANY HEALTH SERVICE	HOW MANY HELPERS WERE SEEN BEFORE ARRIVAL	TIME FROM ONSET TO EFFECTIVE TREATMENT, PER PATIENT	WHICH DOOR ACTUALLY DELIVERS PATIENTS TO THIS SERVICE

The last of the four repays a moment's thought, because a service usually believes it already knows the answer and is usually wrong. A department that runs an outreach programme, a helpline and a referral arrangement with a general hospital will assume its patients arrive in some rough proportion across the three. Counted, the distribution is almost always lopsided, and the lopsidedness is the finding: one route is carrying the service and the others are consuming staff time to deliver very few people. That is an actionable result obtainable from a single field on a form.

Those four quantities, collected on a service's own attenders, answer questions no national dataset can. If the first contact is overwhelmingly one kind of helper, the partnership priority is settled. If the number of steps is large but the interval is short, the problem is misdirection rather than delay. If the interval is long, the question becomes which step it accumulated at. And if one referral source supplies most arrivals, the service knows both where to invest and where it is currently invisible.

The important limitation of these measures is that they describe the people who arrived, which makes them a **clinic-attender sample** and not a population one. They say nothing about those who did not, who are the population of interest and are systematically different. Pathway data from a clinic are therefore a starting point for a district question and not an answer to it, and a service that forgets this will optimise for the population it already has.

38.7 What is not established here, and how to answer without it

This section prints no prevalence for any culture-named presentation, and that is a deliberate position rather than an omission. The figures that circulate for such presentations come from clinic samples of very different kinds, use different criteria, and describe populations that are not comparable; a proportion measured among attenders at one specialist service is not a population

rate and cannot be reported as one. The chapter that owns the clinical material makes the same refusal in print, and this chapter does not supply what that one declined to.

Nor does this section carry an Indian national figure for pathway delay, for the proportion of people whose first contact is a non-medical helper, or for the length of the detour. No such national quantity is held here. The consequence for an answer is straightforward and worth practising: describe the structure of the pathway, name the stages and the levers, say that the national quantities are not established, and point out that a service can generate the local versions itself with four fields on its own assessment form. That is a stronger answer than a remembered percentage, and it is also the more useful one, because a district acts on its own numbers and not on a national average.

TRAP **Answering a question about culture and access with a list of named syndromes.** The roster is memorable, it belongs to another part of the syllabus, and reciting it answers a question about nosology rather than the one that was asked. The question here is about where a population goes, why, what that costs, and what a service does about it. A candidate who names the umbrella once, sends the examiner to the chapter that defines it, and then talks about first contact, detour, recognition and referral has answered the question in front of them.

Culture sets which door opens first; the service decides how far away the second one is, and most of a pathway's length is built from distance, cost and recognition rather than from belief.

39 · Traditional healers, faith-based care and integration with mental health services in India

For a large part of the Indian population the first person consulted about a mental-health problem is not a clinician, and no amount of service planning changes that by wishing. The public-health question that follows is narrow and practical: given that a first contact exists which the health system does not control, what arrangement should the health system have with it. That is a design question about partnership, and it is what this section teaches. The clinical encounter with a patient who has been to a healer, the single measured Indian evaluation of a healing site and what it can and cannot bear, the access-against-acceptability distinction and the clinical posture of collaboration without endorsement are taught in the Consultation-liaison, geriatric and transcultural psychiatry notes. Traditional and indigenous resources considered as an Indian service model in their own right are developed elsewhere in these notes. What belongs here is the architecture of a partnership: what it is for, how referral runs in both directions, what training can honestly attempt, what safeguards are non-negotiable, and how the whole arrangement is held to account.

39.1 The premise, stated without sentimentality

Two propositions have to be held together, and most weak answers drop one of them. The first is that faith-based and traditional healing is not an evidence-based treatment for mental disorder; that is the standing position and this chapter makes no stronger claim for it. The second is that a

great many people consult such practitioners first, that this is a coherent choice given how the problem is understood locally, and that they are frequently nearer, cheaper, more available and more socially acceptable than any psychiatric service.

Held together, those two propositions generate the entire design logic. The case for a partnership is not therapeutic. It rests on **access**, meaning that the traditional healer is a point of contact with a population the service does not otherwise reach, and on **harm reduction**, meaning that whatever happens at that contact can be made safer and shorter. A service that argues for partnership on the grounds that healing works has made a claim it cannot support. A service that refuses partnership on the grounds that healing does not work has declined a route to a population that is going there regardless.

RULE **Partnership with traditional and faith-based practitioners is justified by access and by harm reduction, not by therapeutic efficacy.** Keeping that distinction visible is what allows a clinician to collaborate without endorsing, and it is also what stops the arrangement drifting into an implicit claim that the health system does not intend and cannot defend.

39.2 What a partnership is actually for

Stated as design intentions rather than as demonstrated effects, a partnership is built to achieve a small number of specific things, and naming them is what stops the arrangement becoming a gesture.

- Earlier arrival: a person who is going to reach the health service eventually reaches it sooner, because the healer routes them rather than retaining them.
- Reduced harm at the first contact, meaning that whatever practices carry risk are not applied to people with mental illness.
- Continuity: treatment already started is not interrupted while the person is under a healer's care, which is the commonest concrete damage the health system actually sees.
- Reach in the other direction: the healer becomes a route by which information, and sometimes services, reach a population that does not attend clinics.
- Reduced conflict for the patient, who is otherwise required to choose between two authorities and frequently resolves it by concealing one from the other.

The last of those is underrated and is the one clinicians can verify from their own practice. Where the two systems are understood to be in opposition, patients conceal their use of one, which means the psychiatrist does not know what else is being taken or done, and the healer does not know a medicine has been stopped. A working partnership does not require either side to agree with the other; it requires only that the patient can be honest with both.

39.3 Bidirectional referral, and why the second direction is the hard one

Bidirectional referral is the structural core of the arrangement, and its two halves are not equally difficult. Referral in the obvious direction is easy to describe and hard to obtain: the healer identifies a person who needs medical care and sends them. What that requires is not goodwill but three specific things. A recognisable trigger, meaning a short and unambiguous set

of circumstances in which the healer will refer, rather than a request to exercise clinical judgement. A named destination that is reachable, since a referral into an institution is not a referral. And feedback, because a healer who refers somebody and never learns what happened has been given no reason to refer the next one.

The second direction is the one that decides whether the arrangement survives, and it is the one health services find hardest to accept. A partnership in which the health service receives and the healer only sends is not a partnership; it is a recruitment channel, and it will be recognised as such. What the health service can legitimately send in the other direction is not patients for healing but the things the healer's role actually needs: information about mental illness in usable form, a person to ask when uncertain, acknowledgement of the role the healer plays in the community, and in some arrangements a physical presence at the healing site so that people who will not travel to a clinic can be seen.

PITFALL

Designing the referral in one direction and calling it collaboration. The health service writes a protocol that tells healers when to send people to the clinic, delivers a training session about it, and reports a partnership. Nothing flows the other way, no relationship is built, and the arrangement has no value to the person being asked to change their practice. It lasts as long as the project that funded it, and it teaches the healer community that the health system's interest is extractive.

39.4 Training, and what it can honestly attempt

Healer training is the component most often proposed and most often over-specified. It cannot convert a healer into a health worker, and attempting it fails in both directions: the healer's own practice is what gives them standing in the community, and a curriculum that treats it as an obstacle to be corrected will end the relationship.

What training can attempt is narrow and useful. Recognition of a small number of presentations that need medical attention urgently, described in terms the healer already uses rather than in diagnostic categories. Awareness that certain conditions respond to treatment, which is frequently new information and is the single most persuasive content. The specific harms to avoid, stated as practices rather than as principles. And the mechanics of referral, meaning who to contact, how, and what to say.

Two design points follow from experience with any non-specialist cadre and apply with more force here. Training without continuing contact decays, so a single workshop is an event rather than a programme. And training delivered in a register that positions the trainer above the trainee will not be received, whatever its content, which is why who delivers it matters as much as what it contains.

39.5 Safeguards, which are the floor and not a negotiating position

Everything above is conditional on a **safeguarding floor**, and this is the part of the topic where a clinician's obligation is unambiguous. Practices that restrain, confine, deprive of food or water, inflict physical harm, or hold a person against their will are unacceptable irrespective of the

setting in which they occur or the intention behind them, and a partnership that tolerates them is not a partnership but a complicity. That people with mental illness have been chained and confined in domestic, custodial and faith-based settings in India is documented; how widespread or prolonged the practice has been is a separate question that this chapter attaches no figure to, and none should be quoted without its source, its reference year and its denominator. The statutory prohibition on chaining is legislation and is taught with the rest of the mental health law.

The practical safeguarding content of a partnership is therefore a short list of conditions that are stated at the outset rather than discovered later.

ELEMENT OF THE PARTNERSHIP	WHAT IT REQUIRES TO BE REAL	WHAT IT MUST NEVER BECOME
Bidirectional referral	A recognisable trigger, a named reachable destination, and feedback in both directions	A one-way extraction channel with a protocol attached
Training	Recognition, the treatability message, the harms to avoid, and referral mechanics, with continuing contact	An attempt to convert the healer into a health worker, or a single workshop counted as a programme
Safeguards	An explicit floor on restraint, confinement and deprivation, agreed before the arrangement starts	A matter for later negotiation, or something the health service declines to raise for fear of losing access
Accountability	A named person on each side, and an agreed route for raising a concern	A memorandum with no individual answerable for anything in it
Evaluation	Agreed measures collected from the start, including referral counts and what happened to those referred	A report of activities delivered, with no denominator and no outcome

TRAP **Attributing a partnership policy, a referral protocol or a training scheme for faith healers to an Indian national guideline.** This is the specific error to avoid on this topic, because the sentence writes itself and sounds authoritative. No official Indian guideline position on collaboration or referral between government mental-health services and faith healers is established by anything this chapter has opened. Everything in the design above is a proposal with its reasoning attached, and it should be presented as such. A candidate who says that the design is a reasoned proposal rather than a national policy has answered accurately; one who attributes it to a ministry has invented an instrument.

39.6 Accountability, and the difference between a partner and a supplier

A partnership that nobody is answerable for is a document, and documents do not refer patients. Accountability here means something narrower than it does in a formal service contract, because one side is not employed by anybody and cannot be performance-managed. What it can mean, and what is worth insisting on, is that a named person on each side is the point

of contact, that there is an agreed way of raising a concern in either direction, and that concerns raised are actually answered.

The asymmetry of power in the arrangement deserves stating, because it shapes what accountability can realistically look like. The health service arrives with institutional standing, funding and the implicit authority of the state; the healer arrives with community standing and nothing else that a bureaucracy recognises. In that situation an accountability mechanism designed by the health service alone will read as supervision, and supervision is not what was agreed. The workable version is reciprocal: both sides can raise a concern, both sides answer, and the arrangement can be ended by either.

The measures that matter follow from the purposes named earlier and are few. How many people were referred in each direction over a stated period. What proportion of those referred to the health service actually arrived, which is the number that tests whether the destination was real. What happened to them, in outcome terms the healer can understand and will care about. And whether any safeguarding concern arose, how it was raised and what followed. A programme reporting only workshops held and materials distributed has reported its own activity and nothing about the partnership.

39.7 Making it work in practice

The single grounded statement this chapter can make about how such collaboration actually proceeds comes from an Indian programme in Gujarat that brought faith-based and allopathic practitioners together, and it is about process rather than outcome. Collaboration of this kind **requires trust, rapport-building and open dialogue**, and it is challenging.

That is a modest-sounding finding and it is the most operationally useful thing in the section, because it identifies where the cost of the programme actually falls. The expensive component is not the training curriculum, the protocol or the printed material; it is the time somebody spends building a relationship with people who have no obligation to cooperate and some reason to be wary. That time is invisible in a project plan, is the first thing cut when the timeline tightens, and is what the whole arrangement rests on.

IN INDIA

Because collaboration between faith-based and allopathic practitioners requires trust, rapport-building and open dialogue, an Indian district beginning this work should budget the relationship-building period explicitly and assign it to a named person before any protocol is written. Opening with a memorandum, a training calendar or a referral form inverts the sequence: those instruments are the record of an agreement that already exists, and issuing them first is how a partnership becomes a document nobody honours.

MNEMONIC

Two doors, one floor, one ledger

The shape of a defensible partnership. Referral runs through two doors, not one. A safeguarding floor is agreed before anything else and is not negotiable afterwards. And one ledger records what actually happened, meaning who was referred, in which direction, and what became of them.

39.8 What is not established, and how to answer without it

This is a topic on which confident-sounding claims are unusually easy to produce, so the position is stated plainly.

Not established	Not retrieved	Not quantified here	Process, not outcome
OFFICIAL INDIAN GUIDELINE POSITION ON COLLABORATION WITH FAITH HEALERS	INDIAN DOCUMENT SPECIFYING ALL FIVE PARTNERSHIP ELEMENTS	HOW OFTEN OR HOW LONG THE RELIGIOUS PATHWAY IS TRAVELLED	WHAT THE ONE GROUNDED INDIAN STATEMENT DESCRIBES

No official Indian guideline position on collaboration or referral between government mental-health services and faith healers is available to this chapter. What was sought was exactly such a position, in national operational guidance for mental, neurological and substance-use care; it was not established, and the ownership rule for this material is explicit that partnership design and referral must not be attributed to an Indian guideline. The design set out above is therefore a reasoned proposal, and its authority is the reasoning, not a ministry.

No primary official Indian document or primary Indian evaluation specifying all five partnership elements together, meaning bidirectional referral, healer training, safeguards, accountability and evaluation, was retrieved either. Several routes were attempted. The consequence is that the five-element structure above is offered as a design framework assembled from the requirements of the task, not as a transcription of an existing Indian model, and it should be presented that way.

No national quantification of the religious or traditional pathway is supplied here: how often it is used, and for how long before psychiatric contact, is not something this chapter attaches a number to. The single measured Indian evaluation of a healing site, its denominators, its design limitations and the reasons its result cannot be generalised are taught in the Consultation-liaison, geriatric and transcultural psychiatry notes, and its figure is not reproduced here.

What survives all of that is the durable content, and it is what an examiner is testing. A partnership is justified by access and harm reduction rather than by efficacy. It runs in both directions or it does not run. It rests on a safeguarding floor stated before anything else. It is held by named people rather than by a document. And it is evaluated on referrals and what happened to those referred, not on workshops delivered. Those propositions are defensible without a single figure, which is fortunate, because the figures do not exist.

Collaborate for access and for harm reduction, never on a claim of efficacy; agree the safeguarding floor first, make the referral run both ways, and name the people who are answerable for it.

Consolidate and apply

40 · Worked vignettes

A vignette in this territory is never really about the patient in it; it is about the part of the system that patient is standing in. Stems that carry marks here rarely ask for a diagnosis or a drug. They put an ordinary clinical picture in front of the candidate and then plant one further fact: a distance, a vacant post, a medicine that ran out between visits, a family that consulted somebody else first, a figure quoted at a district meeting. That further fact is the **service variable**, and the value of the stem lies entirely in whether the candidate reasons from it or reads past it into a management plan. This closing section introduces no new programme and no new figure. It rehearses the reasoning assembled across this chapter as a reading skill, working each stem from the service variable outwards to the population it belongs to.

The discipline that earns marks is refusing to answer a service question with a clinical one. A woman who cannot collect her medicine because the facility that stocks it is a bus journey away is not asking which agent suits her; she is asking which tier of the system should be holding her follow-up. Where a stem turns on the statute itself, on the wording of a right, on capacity or consent, or on the procedure that follows a death in custody, it has left this territory: those belong to the Mental health legislation and forensic psychiatry notes and to the Forensic ethics, consent and legal procedure notes. Where it turns on how a prevalence was estimated, or on which design would establish it, it belongs to the Research Methods, Biostatistics and Epidemiology notes. What this chapter supplies is the layer between them: who delivers what, to which population, and how you would know whether they are.

40.1 Four questions locate the service variable

Before deciding what a community psychiatry stem is asking, run the same four questions across it. Between them they cover almost every way an examiner can turn an ordinary clinical picture into a problem about a service, and each opens onto a different part of this chapter.

Level	Offer	Access	Denominator
WHICH TIER IS ASKED TO ACT	WHAT THAT TIER ACTUALLY HOLDS	WHO IS NOT ARRIVING	WHAT THE FIGURE COUNTS

Level asks which tier the question is really being put to: the household and the community worker, the primary care facility and its medical officer, the district hospital, or the specialist institution above them. Offer asks what that tier actually holds, in staff, in supply and in training, as against what the reader assumes it holds. Access asks who is not arriving at all, because a service used by the people who were always going to arrive has changed nothing for the people

who were not. Denominator asks what any figure in the stem counts, who issued it, which period it describes and what it is a proportion of. A well-built stem makes one of the four decisive and leaves the other three as scenery.

RULE **The patient is the setting; the service is the question.** If your answer can be written without naming a tier, a provider, a supply, a pathway or a population, you have answered a question about a person while the stem asked about a system. The clinical picture is usually scenery: it tells you what this person needs and almost nothing about who is meant to provide it. The marks sit in naming the level at which the failure sits, saying what that level does and does not provide, and choosing the move that acts there rather than at the bedside.

40.2 The vignette that resolves on the tier that should act

The commonest stem describes a patient stable for a long time who is not stable now, with the prescription unchanged and nothing new in the history except the travelling. She lives in a village; the clinic that knows her is in the district town; the journey takes most of a working day; somebody has to come with her and lose a day of wages; and the medicine she was given runs out before the next appointment falls due. The reading move is to locate the failure on the ladder of care rather than inside the patient, then ask what the tier nearest her actually holds. A relapse produced by a supply that ran out is a distribution problem; one produced by an appointment the household cannot afford to keep is a geography problem; one produced by a family that stopped believing the treatment was working is a continuity problem. None of the three is answered by changing the drug, and the choice of agent, if the stem really wants one, belongs to the Schizophrenia: management, antipsychotics and adverse effects notes.

The answer that earns marks moves the work downwards and keeps the supervision upwards. Follow-up for a person stable on a settled regimen is the kind of task that can sit with a trained general provider working to a structured protocol, while the specialist tier holds the complicated presentations, the treatment failures and the teaching. That redistribution is what **task-shifting** names, and the intervention guide most often examined for it, with the stepped and collaborative arrangements that make it safe, is developed in this chapter's primary care, task-shifting and stepped-care sections. Where distance is the whole of the problem and a trained provider exists but a specialist does not, a scheduled remote consultation moves the scarce person rather than the sick one, and the conditions attached to doing so sit in this chapter's telepsychiatry section. What a candidate should not write is "refer to the psychiatrist" against a stem whose entire point is that the referral is the thing not happening.

PITFALL

Writing "defaulted follow-up" in the notes without ever asking what the visit costs. The entry travels forward as a statement about the patient's motivation, and every later clinician reads it that way. Ask instead how far the journey is, what the fare and the lost wages come to, who has to accompany her, and whether the medicine lasts until the date you have written on the card. A follow-up plan the household cannot afford to keep has been written for the convenience of the clinic, and it will be read as non-adherence for years afterwards.

40.3 The vignette that resolves on what the programme actually provides

A second family holds the patient constant and moves the question to the district. The stem says the district is covered by the national programme, or that a service was sanctioned for it, and then describes somebody who could not get care there: no clinic ran that month, the psychiatric posts are vacant, the medicines on the district list were not in the store, or the outreach that appears in the annual report has never reached that block. The reading move is to separate **sanctioned coverage** from **functional coverage**. A sanction is an administrative event: an approval, a budget line and a set of posts on paper. Function is a service actually delivered to a population over a stated period, which needs the posts filled, the supply arriving, the staff trained and the activity recorded. The two are counted by different instruments, and the first is far easier to count, which is why it is the figure that travels.

The answer therefore asks what the sanction bought and which of its parts is missing, because the missing part names the level that has to act. Vacant posts are a recruitment and cadre question for the state; empty shelves are a procurement question; untrained staff are a training question; and a clinic that runs while nobody attends is an access question, which is a different vignette again. Naming the missing input is most of the answer, and it is also the honest thing to say about coverage figures in general: what gets published is what the system can count. The components a sanction is meant to buy are set out in this chapter's district-programme section, and the route the money takes to reach them in its financing section.

TRAP **Reading a district count as a service count.** A total of districts approved for a programme is a statement about administrative decisions, not about clinics that opened, and it is quoted as though it described a population served. The same error runs through every coverage figure here: facilities equipped is not facilities functioning, posts sanctioned is not posts filled, and a programme present in a district is not a programme present in every block of it. When a stem offers a coverage figure, the examinable move is to say what it counts before saying what it shows.

40.4 The vignette that resolves on what the figure is a count of

A national number inside a stem is an invitation to interrogate it, not a fact to reason from. The stem quotes something the candidate has heard in a seminar: the proportion of people with a mental disorder who are not in treatment, the calls a national helpline has answered, the density of psychiatrists in the population, the share of the health budget reaching mental health. Then it asks for a conclusion. The reading move is to attach four fields to the figure before using it: who issued it, which period it describes, which population it was measured in, and what it is a proportion of. A figure carrying all four is evidence. A figure carrying none is a rumour with a decimal point in it, and any conclusion drawn from it inherits that weakness.

Two of the traps are structural rather than careless. A cumulative count of contacts with a service, such as the calls a helpline has answered since it opened, counts events and not people: one person in distress may call many times, and no arithmetic on the published total recovers how many individuals it stands for, or how many of them reached care afterwards. And a count with no **valid denominator** cannot be turned into a rate, however inviting the comparison. A national

tally of deaths within an occupational group has no trustworthy count of the people in that occupation to sit beneath it, so a rate built from the two puts an uncertain numerator over an unknown one and reads as precision. The honest sentence names what was counted, says what the denominator would have to be, and stops there. Ask also when the measurement was made, because a household survey describes the period it sampled and not the year in which you are reading it. Where the stem is really asking how such a quantity should be measured, or which design would establish it, that reasoning belongs to the Research Methods, Biostatistics and Epidemiology notes.

The third trap is the one that puts contradictory figures into print. One survey can report the proportion untreated among people with common mental disorders, a second proportion among people with severe mental illness, and a further breakdown for one named condition, and those values are not interchangeable because the populations are not. Each carries a different consequence for which tier should act, since a gap concentrated in severe illness points at the district and specialist tiers while a gap concentrated in common disorders points at primary care. The figures themselves, with the survey that issued them and the populations they describe, sit in this chapter's survey and treatment-gap sections; the disorder-by-disorder breakdown and its epidemiological meaning belong to the Research Methods, Biostatistics and Epidemiology notes.

TRAP **Quoting a survey's figure against a population it did not measure.** One national survey yields several treatment-gap values, one for each population it defined, and the commonest error here is to carry away a single remembered number and attach it to whichever disorder a stem names. The number is then correctly quoted and false at once, which is the hardest kind of error to see, because nothing about the sentence looks wrong. Carry the population with the figure, or carry neither.

40.5 The vignette that resolves on where the pathway to care broke

A third family gives a long history before the first psychiatric contact. A young man has been unwell for years; the family took him first to a healer at a shrine, then to a general practitioner who gave him something for sleep, then to a physician who investigated him thoroughly, and only after all of that to a psychiatric service. The stem is not asking whether the healer did harm, and it is not asking for a formulation of the family's explanatory model at the bedside: that assessment, with the idioms in which distress is described, belongs to the Consultation-liaison, geriatric and transcultural psychiatry notes. It is asking where the **pathway to care** could have been shortened, and by whom.

Read the pathway as a sequence of contacts, each of them a person who could have redirected it. The answer names the point at which recognition failed, the provider standing at that point, and the intervention that acts there: awareness in the community, so that a family reads the problem as one a health service treats; recognition training for whoever the population meets first; a working relationship with the practitioners outside the health system whom a population consults before it consults a clinician; and a service near enough for a referral to be kept. Stigma belongs in this vignette as a service fact rather than an attitude, because it decides who presents,

how late, to whom, and under what description. The strategies, the models of engagement and the evidence for each are set out in this chapter's stigma, literacy and traditional-healing sections.

GOOD TO REMEMBER

Ask every new patient and family two questions before the history proper: who did you see first, and what were you told. The answers give you the route this population actually travels, the contact at which somebody could have redirected it, and the explanation the family is still carrying into your consultation. Collected across a clinic year, those two answers are the closest thing a district service has to a map of its own catchment, and they cost nothing but the asking.

40.6 The vignette that resolves on the population rather than the person

A fourth family describes a death, or a cluster of deaths, and the temptation is to answer with an assessment. A farmer in a district of failed rains and heavy debt has died by suicide using a pesticide kept in the house, and the health administration asks what the service should do. An individual risk assessment is a correct answer to a different question, and the evaluation of risk in the person in front of you belongs to the Somatic-symptom, sexual, impulse-control disorders and suicide notes. What this stem asks is which population lever its facts point at, and they point at several at once: a means kept in every house in the village, an economic shock the health system did not cause and cannot undo, the people the man met before he met any clinician, and a local press that will report the death tomorrow.

The reading move is to separate a **population-level intervention** from a clinical one by asking who receives it. A clinical intervention is delivered to a person already identified; a population intervention is delivered to everyone in a defined group whether or not any individual in it has been identified, which is exactly why it can reach people no service has ever seen. Restricting access to a lethal means, preparing the non-clinicians a distressed person actually meets, agreeing how a death is reported, and acting on the debt and distress that brought the person to the act are all interventions of the second kind. Which one the stem wants is decided by the fact it planted, and this chapter's suicide prevention sections carry what each involves and what the evidence for it is; a vignette answer names the lever and the tier that pulls it, and does not claim a size of effect for it. The same move works on an event that strikes a whole community at once, where an answer built on assessing everybody present individually has answered at a level the event does not permit, and the staged population response sits in this chapter's disaster section.

40.7 The vignette that resolves on what receives the patient

The last family reverses the direction. A man has spent many years in a long-stay institution, his symptoms are settled, the team records that he is fit for discharge, and no family can be traced. Answering with a discharge prescription answers a clinical question that has already been settled; the open question is what receives him. Where people are moved out of institutions faster than the structures they move into are built, the result is **transinstitutionalisation**: the same person held in another closed setting, a custodial one, a shelter or a household with no rehabilitation in it, while the population looks deinstitutionalised because a bed count fell. The rationale, the

outcomes and the lessons drawn from that history are developed in this chapter's deinstitutionalisation section.

So the answer inventories the receiving structures by function: somewhere to live with the support this person needs, something to do in the day, a route towards work where that is realistic, a clinical service that will follow him up without a journey he cannot make, and an entitlement that does not depend on relatives nobody can find. Each is a different service with a different provider, and the provider is often not the state; this chapter's section on Indian models of care sets out the part played by non-governmental organisations, by families and by indigenous resources, and a good answer names who would actually have to be asked. The certification and the statutory entitlement that follow a disability assessment are not this chapter's to state and belong to the Mental health legislation and forensic psychiatry notes.

40.8 Look-alike stems and the answering spine

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

TWO STEMS THAT LOOK ALIKE	THE FACT THAT SEPARATES THEM	THE AXIS IT LANDS ON
A patient relapsing on an unchanged prescription, against one relapsing because the supply ran out	Whether the stem describes a change in the illness or a break in delivery	Level
A district counted as covered, against a district where a clinic actually runs	Whether the figure records an approval or a service delivered over a stated period	Offer
A helpline's total of calls, against the number of people it has reached	Whether the count is of events or of individuals	Denominator
A gap figure for common mental disorders, against one for severe mental illness	Which population the survey measured before any proportion was calculated	Denominator
A long delay before first psychiatric contact, against a family that refused care	Whether anyone met earlier in the sequence recognised the problem and could redirect it	Access
A death that asks for a risk assessment, against one that asks for a population response	Whether the stem names an identified person or a defined group	Level

MNEMONIC

Level, Offer, Access, Denominator

The four service variables spell LOAD, which is what a population puts on a system and what a system is built to carry. **L**evel, the tier the question is being put to. **O**ffer, what that tier actually holds in staff, supply and training. **A**ccess, who is not arriving, and why. **D**enominator, what a figure counts, who issued it, when, and out of what. Each opens a different part of this chapter, and a stem is usually built on exactly one.

With the axis identified, an answer that earns marks follows five moves in order.

- Name the level at which the failure sits, before naming any intervention.
- Say what that level actually provides, and which of its inputs is missing.
- Name who is not reaching the service, and at which contact the pathway broke.
- Attach the issuing body, the period, the population and the denominator to any figure you use, or say plainly that the figure does not carry them.
- Say what you would not do, and why the obvious clinical answer leaves the service variable untouched.

That last move is where public health becomes visible to an examiner. A candidate who answers that the patient should be restarted on medication and reviewed monthly has answered a question about one person, and the answer would read the same in any country and any decade. One who answers that the follow-up belongs at the tier the patient can reach, that the district's problem is a vacant post rather than an absent programme, that the figure quoted describes a different population, and that the people worth worrying about are the ones who never arrived, has answered a question about a service. Every vignette in this territory rewards the second answer.

Find the service variable - level, offer, access or denominator - and answer with the tier that has to act and the population it has to reach, never with the next step in a management plan.

41 · Consolidation and quick review

Nothing in this chapter is about a single consultation. Every section here starts with a population carrying more illness than the service in front of it is reaching, and asks what to do about the difference: a national programme built for districts that had no psychiatric service at all, a survey that counted illness the clinics never saw, a suicide prevention lever aimed at people who will never present, a primary care worker asked to do what no specialist is available to do, and a school, a prison or a camp holding people no clinic will ever see. The individual encounter belongs elsewhere. What belongs here is everybody who never reaches one.

Call that difference the **need-to-service distance**: the gap between what a population needs and what its service actually delivers to it. This closing section teaches no new programme and no new figure, because each is taught in the section that owns it. What remains is to state the

argument those sections were arranged to make, and to convert it into something producible from a blank page with the book shut. The argument is that community psychiatry is the management of that distance, and that the reasoning is the same every time: measure the distance honestly, name who is responsible for closing it, choose the intervention that reaches the people who are not reaching care, and find out whether it did. Questions here are set in exactly that shape, which is why an answer built as a list of programme names scores so poorly against one built as a plan.

41.1 One argument, tested against every population here

Read as a list of territories this material is enormous and disconnected: programmes and policy, surveys and burden, suicide prevention, stigma, rehabilitation, integration into primary care, disasters and special settings, prevention and financing and digital services, culture and traditional healing. Read as one argument it is a single claim tested across all of them. The claim is that each section describes a distance between a population and a service, that the distances differ only in what produced them, and that the same four moves answer every one.

They recur because there are only so many ways a person with a treatable illness fails to end up in treatment. The illness is not recognised as illness. It is recognised, and reaching help is shaming, or is better answered somewhere the family already trusts. It is recognised and acceptable, and there is no service within reach or nobody in it trained to treat. It is affordable to nobody in the household once travel and lost earnings are counted. Contact is made once and never repeated, so what was received was never adequate to the illness. And in some populations the person sits somewhere no service goes at all. Every section here is one of those situations wearing a programme's name, which is why the reasoning does not change as you move through them.

RULE **A service is adequate or inadequate only against a stated population, never on its own description.** The service's own account tells you what it offers; the population tells you what share of the need that offer reaches. An answer that lists a programme's components and stops has described the offer and said nothing about the distance, which is the only thing the programme exists to change.

41.2 Need side or delivery side: the question that opens every answer

Before anything is planned or measured, the distance has to be sorted into one of two kinds, because the two are closed differently and are confused constantly. A **need-side distance** means the population is not arriving: the illness is not recognised, help-seeking is unacceptable, the first person consulted is not a mental health service, or coming costs more than the household will pay. A **delivery-side distance** means the population arrives and the service cannot deliver: no trained worker, no medicine on the shelf, no follow-up, no route onward for the person who does not improve. Both are present in almost every district and neither excludes the other, but at any one moment one of them is binding, and effort spent on the other moves nothing.

The distinction decides the intervention. A need-side distance is answered on the population's own ground: awareness and literacy work, contact with people who have been unwell, respectful arrangements with the practitioners and healers a family consults first, removing cost at the point

of use, and taking the service to the school, the workplace, the prison or the camp. A delivery-side distance is answered inside the health system: training and supervising the workers already posted, moving a bounded task to a worker who can safely do it, securing medicine supply, opening a remote link to a specialist who cannot be posted locally, and arranging the route by which somebody who is not improving is escalated. Candidates who cannot separate the two propose more specialists for a population that was never going to come, and awareness work for a population already waiting outside a clinic with an empty pharmacy.

It also decides what to measure. A population survey measures need and says nothing about what any service did; a service record measures what the service did and says nothing about the people who never came. Only a figure with a service numerator over a population denominator measures the distance itself, which is why **coverage** is the most useful and least available number in this whole area.

41.3 The four moves, and why they run in that order

Once the distance is named and sorted, one sequence organises the answer whichever section it came from. It is a revision scaffold and a planning discipline rather than an official procedure, and it yields in one situation: where a population is in immediate danger, act first and measure afterwards, which is the standing position in the first response to a disaster. Outside that the order matters, because each move is close to meaningless unless the one before it has been made.

- Measure the distance, and measure it honestly. State who issued the figure, what period it refers to, which population it describes, and what it is a proportion of. A figure missing any of the four is not a measurement, it is a quotation. Where it does not exist, say so, and name what would have to be counted to produce it.
- Name who is responsible for closing it. Every distance here sits at a level that can act: a household, a village-level worker, a primary care team, a district unit, a state, a national programme, an employer, a school, an insurer, a court. An intervention proposed without a level that owns it is a suggestion, and it does not survive the meeting it is proposed in.
- Choose the intervention that reaches the people who are not reaching care. Interventions differ far more in whom they touch than in how well they work in those they touch, so selecting on effect alone selects, reliably, whatever works beautifully in the people already coming.
- Find out whether it worked, and design that before it starts. What was the distance before, what is it now, in the same population and against the same denominator. Deciding this at the end leaves only whatever the programme happened to record, which is almost always what it did rather than whom it reached.

Reversing any two produces a failure recognisable in a real service. Choosing the intervention before measuring the distance is how a district acquires the service its planners find most familiar rather than the one its population lacks. Choosing it before naming the owner produces the pilot that works and never scales, because nobody whose job it was ever ran it. And measuring, naming and delivering without ever evaluating is the commonest of the three: it is how a programme

accumulates a shelf of **activity returns** and cannot answer, when finally asked, whether the distance it was built to close has moved.

Distance	Duty	Door	Difference
MEASURED WITH ITS OWN DENOMINATOR	THE LEVEL THAT OWNS CLOSING IT	THE ROUTE THAT REACHES THE ABSENT	THE SAME MEASURE, TAKEN AGAIN

41.4 A frame that names every distance

Because the distances recur and the four moves do not change, one retrieval frame serves the whole chapter. Hold it as a set of prompts rather than as a set of topics: under pressure the failure is rarely that a candidate cannot name a barrier, it is that they answer on the barrier the question made obvious and never scan for the others.

MNEMONIC

Distance, Duty, Door, Difference

The four moves of every answer in this area, in the only order that works. **Distance**: measure the gap between what the population needs and what the service delivers, carrying the body that issued the figure, the period it refers to, the population it describes and the denominator underneath it, and declare it missing when it is missing. **Duty**: name the level that owns closing it, because an intervention with no owner has no route to delivery. **Door**: choose the intervention by whom it reaches, not by how well it works in those already attending. **Difference**: measure the same distance again, in the same population and against the same denominator, and treat a count of what the service did as no answer at all. The same four, asked in the same order, of a national programme, a district unit, a school, a prison or a helpline.

Each distance is taught in full in the section that owns it, with its own evidence and its own figures. The frame below is a retrieval index rather than a teaching table, and it gives every row the same three columns deliberately, because the argument is that the columns do not change.

THE DISTANCE	WHICH SIDE IT SITS ON	WHAT CLOSES IT, AND WHAT WOULD SHOW IT CLOSED
The illness is not recognised as illness	Need side, before any decision to seek help	Awareness and literacy work, and recognition training for whoever is consulted first. Shown by a change in who arrives, not by a count of people who attended a talk
Help is not sought, or is sought elsewhere first	Need side, and decisive where it operates	Contact-based work, change to the rules that disadvantage, and working arrangements with the healers seen before any psychiatrist. Shown by who arrives and how late, not by an attitude score
Care is unaffordable once fee, travel and lost earnings are added	Need side, and the one most often left out of a plan	Cost removed at the point of use, financing that follows the patient, care delivered nearer to people. Shown by uptake among the poorest, not by a scheme existing
No service within reach, or nobody in it trained to treat	Delivery side	Integration into general and primary care, task-sharing with supervision, secure medicine supply, remote links to a specialist. Shown by people treated, not by workers trained
Contact is made once and never repeated	Delivery side, and the most expensive to repair	Follow-up that is somebody's named job, a register of who is due, and a route that escalates when nothing improves. Shown by retention, not by first attendances
The person sits where no service goes: institution, prison, camp, shift, classroom	Both sides at once, which is why these populations are reached last	A service designed for the setting, with the setting's own authority attached. Shown by the population inside the setting, not by the district total

41.5 The measurement layer, and what each figure is for

Numbers are this chapter's connective tissue and its commonest source of error, and the error is almost never arithmetic. Each figure is best held as a **measurement proxy**: a countable thing standing in for something that matters and cannot be counted directly. What a figure is a proxy for tells you what a healthy-looking value fails to prove, which is what is being tested whenever a question hands you a national number.

WHAT IS COUNTED	WHAT IT IS A PROXY FOR	WHAT A GOOD VALUE FAILS TO PROVE
Prevalence in a population survey	How much illness exists, in the population and period surveyed	Anything at all about services. A survey counts people, whether or not one of them ever reached care
Units sanctioned, approved or established	An administrative decision that a service ought to exist	That any of them opened, was staffed or saw a patient. A sanctioned unit is an intention, not a service
Activity: clinics held, workers trained, calls answered	Effort expended, which is the easiest thing any programme can count	That anybody was reached who was not already being reached. Contacts are not people, and one person contacting twice is counted twice
Coverage: those treated over those with the condition	The distance itself, and the only figure measuring it directly	That treatment was adequate. Who counts as treated depends entirely on the adequacy rule applied
Outcome in those treated, such as recovery or remission	Whether the treatment works in the people who received it	That coverage moved. An intervention can produce excellent outcomes in everyone it reaches and leave the distance where it was
Money allocated, released or spent	Stated priority, and separately the capacity to convert money into service	That the service exists. These are three different figures, and only the third attaches to something that happened

Held this way the numbers stop being a list to memorise and become a set of questions, and the drill worth doing is the reverse one. Take a figure of the kind a question will hand you and, before commenting on its size, produce its population, its period, its issuing body and its denominator. Where one of those cannot be supplied, naming the missing field is a stronger answer than a confident number.

41.6 When closing the distance is not the aim

Two situations sit outside the argument as stated, and applying the frame mechanically to either is wrong in a way that is easy to hear. The first is the **population lever**, the intervention that never touches care-seeking at all. Restricting access to a lethal method, changing how a death is reported, designing work so that it does less harm, amending a rule that disadvantages people with mental illness: none of these depends on anybody recognising an illness, deciding to seek help or reaching a service. They act on everybody at once, they are usually delivered by somebody outside health, and their effect appears in no service record. The question to ask of them is not who was reached but whether the exposure changed, and the level responsible is often a ministry, a regulator or an employer rather than a health service.

The second is the distance that cannot be measured with the counts available. Some of the most quoted figures in this field have no valid denominator: a cumulative total of contacts with a service is not a count of people, and a count of deaths in an occupational group has no matched population figure to sit over. The honest position is that the numerator is real, the denominator does not exist in matchable form, and no rate may be built from the pair. Saying so is not a failure to answer. Building the rate anyway is the failure, and it is the one that gets quoted onward by everybody who hears it.

41.7 The two ways candidates lose these marks

Two failures account for most of what is lost here. The first is **catalogue answering**: reciting programmes, components and structures without once saying what distance any of them addresses or how anyone would know it had closed. It reads as prepared, it fills a page, and it answers an administrative question nobody asked. The second is quoting a national figure without the population it belongs to, which is worse than it looks, because a correctly transcribed number attached to the wrong group is a false statement and no arithmetic check will catch it.

TRAP **Answering a service question with more specialists.** Asked how to improve care for a district, a population or a special setting, candidates reach almost reflexively for more psychiatrists and more clinics, which is a delivery-side answer given without checking which side the distance is on. Where the population is not arriving, capacity added at the clinic door changes nothing and shows up in no measure. An answer that names the side first, then the level responsible, then the intervention, then the measure that would show it worked, has earned most of what is available whatever it eventually proposes.

PITFALL

Reading a busy service as a reached population. A unit that is full, whose staff are exhausted and whose returns look excellent, tells you what happens to the people who arrived and nothing whatever about those who did not. The two move independently, and the busiest services are often the ones with the least idea of their own catchment. The habit worth keeping is to ask, of any service you work in, what share of the people in its area with the condition it treats has ever reached it. Where nobody can answer, that is the finding.

41.8 Where this chapter stops, and what to carry in

An answer is judged as much by what it declines to teach as by what it covers. Consolidating this material means consolidating its edges, because several neighbouring chapters treat the same programmes, populations and deaths from a different angle. Each of the following is worth rehearsing as a handover: named where it touches the question, passed to its owner in a sentence, and left there.

- The provisions of mental health legislation, the advance directive, the nominated representative, the review machinery and the decriminalisation of attempted suicide belong to the Mental health legislation and forensic psychiatry notes. What is kept here is the service consequence: who arrives, what the service records, and what capacity the change assumes.

- Professional conduct, consent, capacity at the bedside, custodial and court procedure and the statutory arrangements for disaster response belong to the Forensic ethics, consent and legal procedure notes.
- Incidence, prevalence, study design, sampling and the derivation of any rate belong to the Research Methods, Biostatistics and Epidemiology notes. This chapter uses those measures as instruments and never derives them.
- The assessment and management of suicide risk in the person in front of you belongs to the Somatic-symptom, sexual, impulse-control disorders and suicide notes. What is kept here is the population lever.
- Cultural formulation at the bedside, idioms of distress, the named culture-bound presentations and the assessment of somebody who has consulted a faith healer belong to the Consultation-liaison, geriatric and transcultural psychiatry notes. What is kept here is the pathway and the partnership.

GOOD TO REMEMBER

Before writing a line of any service proposal, put the four moves down as four headings and fill them in that order. The discipline shows in the fourth: deciding what measurement would demonstrate a change usually forces a change in the intervention itself, because attractive proposals often turn out to move nothing anybody can count. Choosing the measure last, once the plan is fixed, is how a service ends up with years of returns and no evidence.

Carried into the examination, all of it collapses to a sequence that survives when the detail fails. Name the population and the distance. Say which side is binding. Name the level responsible for closing it. Choose the intervention on whom it reaches. State the measure that would show the distance had moved, and say plainly where that measure does not exist. Consolidation buys not more facts but faster access to the ones already learned, so that under pressure the answer arrives as a method rather than as a half-remembered list of programme names.

Community psychiatry is the management of the distance between what a population needs and what its service delivers: measure the distance with its denominator, name the level responsible for closing it, choose the intervention by whom it reaches rather than by how well it works in those already attending, and then measure the same distance again.

PART XI

XI

Medical and Systemic-Disease Psychiatry

Psychiatry at the medical interface: the heart, the lungs, the failing organ, the transplant and the cancer, and the discipline of routing what this chapter does not own.

01 Cardiac, respiratory and critical-illness psychiatry

02 Renal, transplant and perioperative psychiatry

03 Oncology, palliative and paraneoplastic psychiatry

04 Autoimmune and inflammatory systemic disease, and burn injury

05 Consolidate and apply

VISUAL 11.1



The organ interface changes psychiatric decisions

- Disease physiology
- Psychiatric symptoms
- Psychotropic exposure
- Organ adjustment
- Liaison formulation
- Monitoring

Medical psychiatry asks how organ disease, psychiatric symptoms and treatment exposure alter one another and what must be monitored.

Source lock: day-40.final.html#sec-t01

VISUAL 11.2



Critical illness has a long psychiatric shadow

- Delirium
- Anxiety and PTSD
- Physical recovery
- Cognition
- Family experience
- Follow-up

Post-intensive-care outcomes include psychological, cognitive and physical problems in patients and families and require longitudinal follow-up.

Source lock: day-40.final.html#sec-t03

VISUAL 11.3



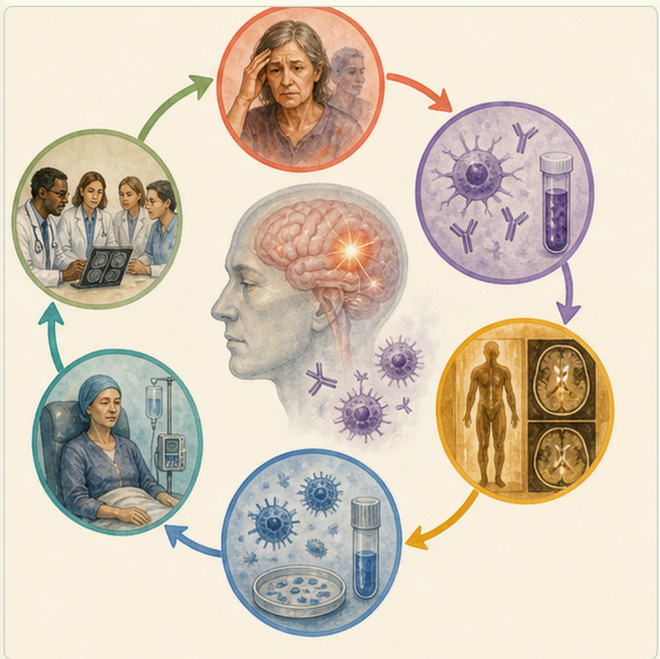
Palliative psychiatry separates overlapping distress

- Depression
- Demoralisation
- Delirium
- Existential distress
- Symptom care
- Shared goals

Palliative psychiatry distinguishes depression, demoralisation, delirium and existential distress while aligning symptom care with the patient's goals.

Source lock: day-40.final.html#sec-t06

VISUAL 11.4



Systemic disease can signal brain inflammation

- Neuropsychiatric red flags
- Immune evidence
- Malignancy search
- Exclude infection
- Immunotherapy
- Coordinate specialties

Autoimmune, inflammatory and paraneoplastic syndromes are approached through red flags, staged exclusion, immune evidence and coordinated treatment.

Source lock: day-40.final.html#sec-t16

PAPER IV

Medical and systemic-disease psychiatry

Four bands . one complete notes document

Cardiac, respiratory and critical-illness psychiatry

1 · The organ interface: what this chapter owns, and what it deliberately points to

At the interface with medicine the psychiatrist is rarely asked what the diagnosis is, and is almost always asked which of three things is producing the picture. The medical illness may be generating the psychiatric symptoms directly. The treatment given for that illness may be generating them. Or the patient may have a psychiatric disorder of their own, arriving in a body that happens to be unwell. Those three answers oblige three different actions, and this chapter is built to make the choice between them quicker and more defensible. This orientation installs that **triage**, adds the fourth question the triage forces, which is who teaches a construct once it has been placed, and draws the map of what follows.

One feature of this chapter should be stated at the outset rather than discovered halfway through. It is deliberately short for the number of constructs it names, because several of them are named here and taught elsewhere in the series. That is a design decision with a defence, and the defence is set out below in full, because a reader who does not understand why a chapter points will read the pointing as a shortfall.

1.1 The three-way question, and what each answer obliges you to do

The first limb is the illness acting on the brain. A systemic disease can produce psychiatric symptoms through mechanisms that require no psychological step at all: inflammation reaching the central nervous system, an immune response raised against a tumour that also recognises neural tissue, hypoxia and carbon dioxide retention, the accumulated products of failing kidneys, endocrine derangement, small-vessel injury, a surge of circulating catecholamines. When this limb is the answer, the psychiatric task is largely to establish it and to hand the patient back for treatment of the disease, because no psychiatric intervention will hold while the process continues.

The second limb is the treatment acting on the brain. Modern medical treatment is potent and it is neuroactive, and the psychiatric consequences of it are not rare accidents but predictable properties of certain agents and certain procedures. Immunosuppression, cytotoxic treatment, interferon, corticosteroids, anaesthesia and surgery itself, oxygen and ventilatory support and implanted devices all change the mental state, each with its own syndrome and its own interval. When this limb is the answer, the decisive evidence usually sits in the drug chart and the procedure record rather than in the mental state examination, and the action is a conversation with the treating team about the exposure.

The third limb is a psychiatric disorder in its own right, occurring in a person who is medically unwell. This limb has two internal shapes that behave alike for management purposes. The

disorder may be one the patient would have developed anyway, whose timing merely coincides with the admission. Or it may be a response to the illness, to what the illness has taken, and to what it will take next. Both need treating as psychiatric conditions, and both are vulnerable to the same argument, which is that a person in this situation would reasonably feel that way. When this limb is the answer, the psychiatric task is the largest and the least likely to be delegated.

Two properties of the triage matter more than the categories themselves. The limbs are not exclusive, and in a medically unwell patient more than one is usually operating at once, so the working question is not which limb is true but which limb is doing the most damage now and which, if missed, is dangerous. And the limbs are answered by different evidence: the first by the disease course and its physiology, the second by exposure and timing, the third by the patient's history, baseline and account of themselves.

Illness	**Treatment**	**Independent disorder**	**Owner**
READ FROM THE DISEASE COURSE	READ FROM EXPOSURE AND TIMING	READ FROM BASELINE AND HISTORY	WHO TEACHES WHAT YOU HAND ON

RULE **State which limb you are in, and say what evidence put you there.** An answer at this interface that names a psychiatric syndrome without placing it against illness, treatment or independent disorder has skipped the step the chapter examines. The examiner's follow-up is almost always the same question in different words: what else could be producing this, and what would separate them?

1.2 The examination technique the triage runs on

The triage is decided by history far more often than by investigation, and by a small number of manoeuvres that are cheap and frequently skipped. Learning them as a fixed order is worth more than learning any single association in this chapter, because the order transfers to every section in it.

Establish the baseline first, from someone other than the patient. What was this person like a month ago, at work, at home, in conversation? A medically unwell patient is observed by staff who never met the well version, so the comparison that would make the change visible does not exist on the ward unless a psychiatrist goes and gets it. Without a baseline, a description of the present state is the only thing available, and a present state cannot establish a change.

Then establish **tempo**. Every psychiatric illness has a characteristic time signature, and the useful reading is the mismatch: a presentation moving faster than the condition it resembles, or continuing to move while adequate treatment is being given, or following a slope rather than an episode. Tempo costs nothing but a careful timeline with corroboration, and it is the reading most often left vague.

Then lay two timelines side by side. The first is the medical course, with its events dated: the acute episode, the operation, the transfer, the deterioration. The second is the drug chart, read for start dates rather than for the current list. Laying the psychiatric change against those two timelines is the single highest-yield act at this interface, and it is what separates the first limb

from the second. A change that begins with the illness and tracks its severity points one way. A change that begins after an agent was introduced, and in the interval characteristic of that agent, points the other.

Ask what was stopped as well as what was started. An admission removes access to alcohol, to nicotine and to whatever else the patient used at home, at a predictable hour, and it frequently interrupts the patient's own long-term psychotropic treatment on the ward round of a team that does not prescribe it. A withdrawal syndrome or an abrupt discontinuation, declaring itself on a medical ward, is read by people looking for a medical explanation as a complication of the illness.

Attend to consciousness and attention before anything else in the mental state. A fluctuating disturbance of attention and awareness changes the entire answer and moves the problem into territory this chapter does not teach, and the whole of that assessment, including any bedside instrument, rating scale or drug choice for it, belongs to the stroke, traumatic brain injury, delirium and focal brain syndromes notes. Recognising the shape and handing it on is the examinable act here.

TRAP **Settling for understandable.** The commonest error at this interface is not a wrong diagnosis but an abandoned one: the low mood is judged a reasonable response to a serious illness, and the assessment stops. Comprehensibility is not a diagnostic exclusion and would not be accepted for any other comorbidity. It costs more than usual here, because the psychiatric state is often what removes the patient from the treatment, the rehabilitation and the adherence that would change the medical trajectory.

1.3 Why a chapter that points beats a chapter that teaches everything thinly

The alternative to routing a topic away is not depth, it is a thin second copy, and a thin copy is worse than an absence because it does not feel like one. A paragraph read on a topic taught properly elsewhere leaves the reader recognising every term, able to defend none of them, and unable to tell which of their own facts are the thin ones. A half-learned fact sits in memory looking exactly like a fact learned well. It survives the first question and fails the follow-up, which is where the marks are, and it fails at the bedside for the same reason.

An explicit omission behaves differently. A reader who has been told that a construct is taught in a named chapter knows precisely what they do not yet have, and can act on it. A reader who met a thin paragraph knows nothing of the kind. That difference is also what makes routing checkable rather than merely tidy: whether a chapter mentioned a topic is a weak test that a chapter can pass while leaving its reader unable to answer anything, whereas whether the reader can name the owner is a test with a right answer.

There is a second argument from the reader's side, and in the last weeks before an examination it is the stronger one. Revision time is fixed, and a chapter that duplicates an adjacent chapter spends that time twice on one topic and not at all on the topics only it can teach. The material this chapter alone carries, which is the systemic and inflammatory disease that presents psychiatrically before anyone knows it is present, is exactly the material that goes unrevised when a chapter pads itself out to look complete. Shortness here buys depth where depth is scarce.

A third argument concerns the text rather than the reader, and this chapter's routing section makes it in full: two accounts of one topic, written by different hands, drift apart, and one owner with one account is the only arrangement in which two pages of the same book cannot disagree.

The fourth argument is the one that matters clinically, and it is why the routing question is not bookkeeping. Liaison work is **routing**. A referral arrives from another team, written by someone unsure what to ask, and the psychiatrist's first act is to decide whose question it is: the treating team's, the psychiatrist's, or a third specialty's altogether. A resident who has practised deciding where a construct belongs has been practising the actual skill, and one who has memorised a paragraph about each has not. The map that follows is therefore content, not furniture.

RULE **A declared gap is a correct answer, and a supplied specific is not.** Where this chapter does not hold a figure, a threshold or a management sequence, it says so and names the kind of source such a statement would rest on. In an examination the equivalent move is to state the direction confidently, state that the magnitude is not something you can quote, and go on. That answer is stronger than a number produced from recollection, which is unverifiable and frequently wrong.

The map of routed constructs is drawn once, on the plate below, and the routing section that follows gives each one an exact address and a single boundary sentence stating what this chapter withholds about it.

1.4 The map of the chapter, and the question each section answers

The chapter is organised into four bands. The first covers cardiac, respiratory and critical-illness psychiatry. The second covers renal, transplant and perioperative psychiatry. The third covers oncology, palliative and paraneoplastic psychiatry. The fourth covers autoimmune, rheumatological and inflammatory systemic disease. A section carrying the psychiatric care of the severely burned patient sits alongside them, and a routing section carries the constructs taught elsewhere. Each section below is named by the question it answers, because a question is easier to hold in mind than a title and it is what an examiner is actually setting.

SECTION	THE QUESTION IT ANSWERS	LIMB OF THE TRIAGE IT MOSTLY TRAINS
Heart failure, arrhythmia and cardiac devices	Which of the somatic symptoms belong to the failing circulation, how do rhythm presentations separate from panic, and what does an implanted device do to a person and a household?	Illness, and independent disorder
Obstructive airways disease, asthma and chronic respiratory failure	Is this breathlessness or is it panic, and which behavioural levers change the trajectory?	Illness, and independent disorder
The stress-heart interface	Can an acute psychological event produce a structural cardiac lesion, and why is the presentation mistaken for a coronary event?	The arrow reversed, running from the psychological event to the organ
Transplant psychiatry and immunosuppressant neurotoxicity	What must a pre-transplant psychiatric opinion contain, and which new neuropsychiatric states belong to the drugs that keep the graft?	Treatment
Perioperative psychiatry and postoperative neurocognition	In what order do the perioperative questions arrive, and which cognitive construct is being named after the operation?	Treatment
Cancer-drug neuropsychiatry	Which agent, which syndrome, and over what interval does that syndrome appear and get looked for?	Treatment
Paraneoplastic neuropsychiatric syndromes	When should a psychiatric presentation start a directed search for a malignancy nobody has found yet?	Illness, before the illness is known
Neuropsychiatric lupus	Does the syndrome in front of you belong to the disease at all, and how is that attribution reasoned when no test settles it?	Illness, with attribution as the whole difficulty
Neurosarcoidosis, cerebral vasculitis and antiphospholipid syndrome	How do you answer on a rare systemic cause without inventing a phenotype for it?	Illness, and evidential discipline
Coeliac disease, inflammatory bowel disease and the gut-brain axis	What do the association and bidirectionality findings actually state, and what will their design bear?	Illness, and independent disorder

SECTION	THE QUESTION IT ANSWERS	LIMB OF THE TRIAGE IT MOSTLY TRAINS
Burns psychiatry	What is the psychiatric task at each stage of the burn timeline, and what question does the injury itself raise about what preceded it?	Independent disorder, and response to injury
Organ-interface routing	Who teaches the constructs this chapter names and does not teach, and what exactly is withheld here?	The fourth question

Read the third column across and the lens becomes visible in the shape of the chapter. The sections on paraneoplastic syndromes, lupus and the rare systemic autoimmune causes are extended training in the first limb, and specifically in its hardest form, where the illness has not yet been identified and the psychiatrist is the person who must raise the possibility. The sections on transplantation, the perioperative period and cancer drugs are training in the second limb, where the answer is in the exposure and the interval rather than in the mental state. The cardiac, respiratory, bowel-disease and burns sections spend most of their length on the third limb, on psychiatric disorder in a person who is genuinely and often permanently unwell, and on the reasoning that keeps it treatable rather than merely understandable.

One section sits outside the scheme on purpose. The stress-heart interface runs the arrow the other way, from an acute psychological event to a demonstrable cardiac lesion, and it is in the chapter because it is the clearest available demonstration that the traffic at this interface is not one-directional. Holding that section next to the triage is useful precisely because it is the exception.

1.5 What this chapter does not teach, and where each of those lives

Five boundaries are enforced throughout, and knowing them in advance saves a reader from searching this chapter for material it has deliberately excluded.

- Prescribing at the interface, meaning organ-adjusted dosing, interaction rankings and the choice of agent in cardiac, hepatic, renal or transplant disease, belongs to the psychopharmacology notes. This chapter teaches everything that precedes a prescription and adds no second version of the prescription itself.
- The assessment and management of the acute confusional state belongs to the stroke, traumatic brain injury, delirium and focal brain syndromes notes, including every bedside instrument, rating scale and drug choice for it. Several sections here reach that boundary, at the end of life, after an operation, in critical illness and in the early burn injury, and each marks the hand-off rather than crossing it.
- Capacity belongs to the mental health legislation notes. This chapter says only that capacity is decision-specific and time-specific, which is enough to reason with, and it states no statutory test.

- The generic immune-brain framework, the workup of autoimmune encephalitis and the neuropsychiatry of interferon-beta belong to the neurology and psychiatry interface notes. What this chapter owns on the immune side is the disease-specific presentation and the decision to suspect, not the general model and not the investigative protocol.
- The constructs defined by a clinical setting rather than by an organ, being the coronary care unit, the intensive care unit, the dialysis unit, the oncology ward and the palliative care unit, belong to the consultation-liaison, geriatric and transcultural psychiatry notes, and the psychiatric disturbance that follows a corticosteroid belongs to the endocrine and metabolic psychiatry notes. Those two owners take the constructs this chapter routes, and the routing section names each one with its boundary.

None of this is a shortfall being managed politely. What those constructs share is a referral pathway and a liaison relationship, not a pathophysiology, so they belong with the chapter that teaches liaison work. The one that follows a drug follows the drug. The classification is working, and stating it out loud is the difference between a boundary and an omission.

1.6 How to use this chapter in the weeks before the examination

Read the routing plate before anything else, so that the constructs this chapter will not teach are settled and stop feeling missing. Then read each section for its discriminators rather than for its lists, because every section here is built around a small number of separations that decide the answer, and the lists around them are recognisable at the paper and not defensible under questioning.

Practise the triage out loud on whichever medical patient you happen to be asked about, and practise it as a four-step answer rather than as three categories. The four cues below are a teaching scaffold for that sequence rather than a validated instrument, and the sequence itself is the answer structure this whole chapter installs. It works on a construct you have never met, which is more than can be said for a memorised list.

MNEMONIC

Limb, Evidence, Shift, Owner

Limb: which of the three the picture mostly sits in, the illness acting on the brain, the treatment acting on the brain, or a psychiatric disorder in its own right in a person who is unwell. **Evidence**: what placed it there, which is the disease course and its physiology for the first limb, exposure and timing for the second, and the patient's baseline, history and account of themselves for the third. **Shift**: what would move you to another limb, said out loud, because the limbs are not exclusive and the working question is which one is doing the most damage now. **Owner**: which chapter takes whatever you have just handed on, named rather than rebuilt from memory.

IN INDIA

No India-specific epidemiology is held for this orientation and none has been invented. The Indian point here is one of service structure. In most Indian general hospitals the psychiatrist arrives at the medical bedside on request, usually late, usually to a referral question written by a team under pressure, and often without the previous notes, the collateral history or the drug chart's start dates to hand. Every manoeuvre in the triage above depends on information the referring team holds and the referral letter rarely contains, so the practical skill is asking for it explicitly and early rather than working from the mental state alone. That is also the argument for reading the consultation-liaison material before a liaison posting begins.

GOOD TO REMEMBER

Three questions at the bedside and one afterwards. Is the disease doing this? Is the treatment doing this? Is this a psychiatric disorder in a person who is unwell? And then, before writing: whose chapter answers the part I am handing on?

This chapter teaches a triage and a map, and it teaches them together, because at the medicine interface deciding where a problem belongs is most of deciding what to do about it.

2 · Depression and anxiety after myocardial infarction and in coronary artery disease: prognosis and safe antidepressant choice

Six of the topics in this chapter are addresses rather than lessons, and saying so plainly is the most useful thing this section can do. Depression and anxiety after myocardial infarction and in coronary artery disease, intensive care psychiatry and the psychological sequelae of critical illness, psychiatric assessment in end-stage renal disease and on dialysis, depression and demoralisation in cancer, palliative care psychiatry, and steroid-induced psychiatric disturbance are each named in this chapter and taught in none of it. Each has an owner elsewhere in the series. This section hands the reader that owner, states what is withheld here, and explains why withholding it is honest rather than lazy.

2.1 Why a routed chapter is more honest than a complete-looking one

The alternative to routing is not depth. It is another thin copy. A chapter that gives every adjacent topic a paragraph produces a reader who recognises every term on the paper, can defend none, and cannot tell which of their facts are the thin ones. That is the harm, and it is not obvious: a half-taught fact does not feel half-taught. It sits in memory looking like a fact learned properly, and it fails at the point of use rather than the point of reading.

The second harm is duplication, and it is invisible while writing. Two accounts of one topic, written by different hands against different sources, do not stay identical. They drift, and a reader who has met both cannot tell which is authoritative. One owner and one account is the only arrangement in which two pages of the same book cannot disagree.

The third point is about how completeness gets checked, which is why routing must be declared out loud rather than practised quietly. Asking whether a chapter mentions a topic is an existence check, and a chapter can pass it while leaving the reader unable to answer anything. Asking whether the reader can name the owner is the better test, because it is checkable. A **routed construct** turns an implicit omission into an explicit one, and an explicit omission can be acted on. Silence cannot.

A fourth reason is particular to this chapter. Five of the routed topics are not defined by an organ at all. They are defined by a setting: the coronary care unit, the intensive care unit, the dialysis unit, the oncology ward, the palliative care unit. What they share is a referral pathway and a liaison relationship, not a pathophysiology, so they belong together in the chapter that teaches liaison work. The remaining topic is defined by a drug and follows the drug. Routing here is the classification working correctly, not a shortfall being managed politely.

RULE **Each routed construct gets one exact cross-reference and one boundary sentence here, and nothing else.** This section supplies no incidence, no prognosis, no management sequence and no dose for any of them, and that restriction is its design, not its limitation. A sentence here that supplied one would be the second copy this arrangement exists to prevent.

2.2 The routing plate, and how to read it

The plate below is this section's content, and each row holds three things.

- The construct, worded as the examination will name it.
- The chapter that owns it, which is what the reader is meant to leave with.
- The single sentence stating what this chapter will not say about it, so that the omission is explicit rather than silent.

The rows are addresses with a warning label attached; they will not survive being revised from.

The owning-chapter column is repetitive, because most of these are liaison presentations pointing at the same place. Some rows carry a split, because a construct can have its assessment owned by one chapter and a single limb owned by another, and the plate says so.

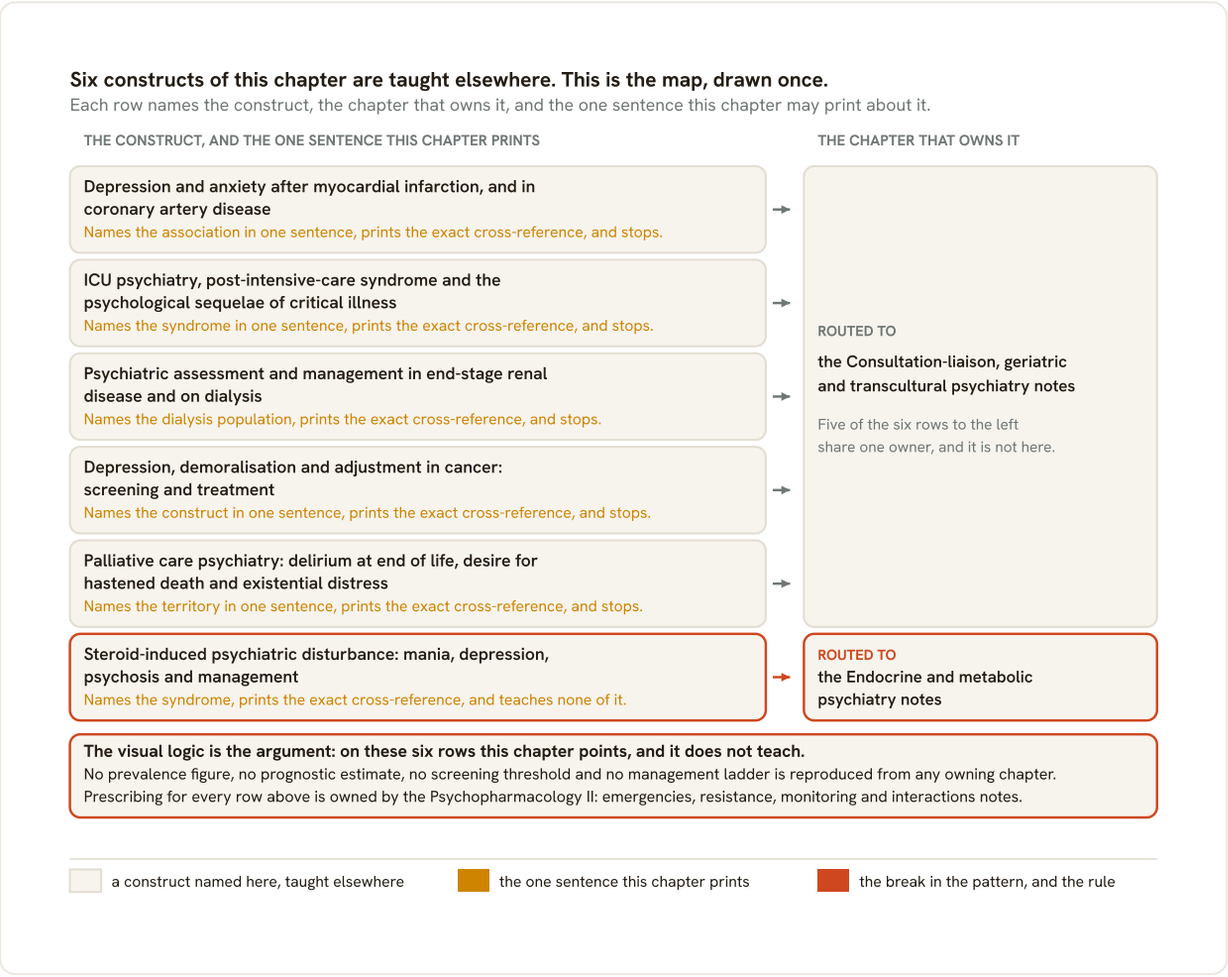


Fig 40.1 The routing map, which is the whole visible content of the organ-interface routing section. Six of this chapter's constructs have owners in earlier chapters, and drawing that once is what makes a short chapter honest rather than evasive. Read the shape before the words: five rows converge on a single destination panel, so the reader sees at a glance that consultation-liaison psychiatry, not this chapter, holds post-infarction mood, critical-illness sequelae, dialysis, cancer and palliative care. The sixth row breaks away because steroid psychiatry belongs to the endocrine and metabolic chapter, and that break is the only thing the second colour means. What each row prints in amber is the entire permitted output: a naming sentence, the exact cross-reference, and a full stop. From memory you should be able to say, for any of the six, which chapter owns it and that this one adds no figure, no threshold and no ladder of its own.

2.3 What this section cannot establish, and how to answer without it

Two things a fuller account would supply are declared here as gaps, because a **declared gap** is a correct answer and an invented figure is not.

The first concerns anxiety rather than depression after infarction. A statement of its prognostic effect would rest on a meta-analysis in the cardiology literature, identified for this series but not opened in quotable form. The consequence is stated rather than concealed: this chapter attaches no prognostic figure to post-infarct anxiety, and a candidate should attach none from memory. What can still be said is that anxiety and depression after infarction are separate constructs on separate evidence, and that an answer treating them as interchangeable makes a claim it cannot support.

The second concerns the desire for hastened death in advanced illness. No current guideline recommendation, and no primary intervention study, could be opened for this series with text usable on its own terms on assessing or managing it; such a claim would rest on an oncology

clinical-practice guideline group, or on investigators of primary palliative-care psychotherapy trials. The absence of a quotable recommendation licenses no inaction: a statement of that kind is met as a communication carrying meaning rather than a symptom to be corrected, and decision-making capacity, which is decision-specific and time-specific, is doctrine owned by the Mental health legislation and forensic psychiatry notes. What is withheld here is a protocol, not the obligation to respond.

2.4 The rows in words, with the boundary stated for each

Depression and anxiety after myocardial infarction and in coronary artery disease is owned by the Consultation-liaison, geriatric and transcultural psychiatry notes, and the choice of antidepressant in established cardiac disease is owned by the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes. This chapter states neither the size of the prognostic effect nor the name of any preferred agent.

3 · ICU psychiatry, post-intensive-care syndrome and psychological sequelae of critical illness

Intensive care psychiatry and the psychological sequelae of critical illness are owned by the Consultation-liaison, geriatric and transcultural psychiatry notes, and delirium in the critically ill patient is owned by the Stroke, TBI, delirium and focal brain syndromes notes. The route follows from the shape of the construct, since post-intensive care syndrome encompasses physical, cognitive, and psychological impairments that persist in patients after discharge from an intensive care unit. Only one of those domains is psychiatric, so a chapter organised by organ could own only a slice of it. Its assessment, its course and its treatment appear here in no form.

4 · Psychiatric assessment and management in end-stage renal disease and on dialysis

Psychiatric assessment and management in end-stage renal disease and on dialysis is owned by the Consultation-liaison, geriatric and transcultural psychiatry notes, and dose adjustment in renal impairment is owned by the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes. The routing follows the service design rather than the diagnosis, since psychological support and interventions should be delivered within the kidney care setting. That makes it a liaison topic before a nephrology one, and no screening approach, treatment sequence or dosing rule appears here.

5 · Depression, demoralisation and adjustment in cancer: screening and treatment

Depression, demoralisation and adjustment in cancer is owned by the Consultation-liaison, geriatric and transcultural psychiatry notes. That chapter carries the distinction between demoralisation and depressive disorder, and the treatment sequence; this chapter carries neither, so a candidate revising how demoralisation differs from depression should not look here. A screening question or brief instrument for depression in advanced disease is stated by neither chapter, and that is a gap in the series rather than a section you have failed to find, so supply none from memory.

6 · Palliative care psychiatry: delirium at end of life, desire for hastened death and existential distress

Palliative care psychiatry is owned by the Consultation-liaison, geriatric and transcultural psychiatry notes, with delirium as a syndrome owned by the Stroke, TBI, delirium and focal brain syndromes notes and the doctrine of capacity owned by the Mental health legislation and forensic psychiatry notes. This chapter says only that capacity is decision-specific and time-specific, and states no approach to sedation, existential distress or a request for hastened death. What changes about delirium once treatment intent has changed, including how far to investigate and how antipsychotic choice differs in the final phase, is stated by neither owner and is a gap in the series rather than a section you have failed to find.

7 · Steroid-induced psychiatric disturbance: mania, depression, psychosis and management

Steroid-induced psychiatric disturbance is owned by the Endocrine and metabolic psychiatry notes. It is the row that follows a drug rather than a setting, so its dose relationship, its onset window and its management sequence all sit in that chapter and none is stated here. The chemotherapy and interferon section of this chapter carries no steroid limb either, by owner ruling, so there is exactly one place to look.

TRAP **Revising from the routing line as though it were an abstract.** A cross-reference is an address, written to be too short to revise from. The failure mode is a candidate who says depression after infarction matters to prognosis but cannot say what it does to it, or who knows corticosteroids cause psychiatric disturbance but cannot say when. Both collapse on the follow-up.

IN INDIA

No India-specific row is held for this section, and none has been invented. The Indian point is one of study order rather than epidemiology. Each of these constructs is by definition a liaison presentation: the patient is admitted under another team, and the psychiatrist arrives on request, often late and with a referral question written by somebody unsure what to ask. The chapters that own them also carry the referral, capacity and communication material, so read them before a liaison posting starts.

GOOD TO REMEMBER

At the bedside the question is not what you remember but which chapter answers the referral in front of you. Cardiac ward, intensive care, dialysis unit, oncology ward and palliative care unit all route to the liaison chapter. A psychiatric disturbance that began after a corticosteroid was started routes to the endocrine and metabolic chapter. Knowing that split in advance beats a half-remembered paragraph about any of them.

These constructs are named here, routed here and taught nowhere in this chapter, so that each is taught once, in full, where it is owned.

8 · Psychiatric aspects of heart failure, arrhythmia and cardiac devices

The failing heart produces most of the somatic symptoms a depression scale counts, so cardiac psychiatry is decided by history and by stage of illness rather than by a symptom total. Drug treatment is deliberately absent: prognosis and antidepressant choice in cardiac disease are taught with depression and anxiety after myocardial infarction and in coronary artery disease, and organ-adjusted dosing with the psychopharmacology notes. What is examinable here precedes any prescription.

8.1 The overlap that defeats a symptom count

Advanced heart failure and a depressive episode share almost the entire somatic side of the syndrome. Breathlessness on exertion, fatigue, disturbed sleep, poor appetite, weight change, psychomotor slowing and loss of sexual interest each follow from a failing circulation, from the treatment burden it imposes, or from a mood disorder, and none of them separates the possibilities alone. A scale that counts them inflates where accuracy matters most; a clinician who discounts them all misses the depression that is present.

Weight the **cognitive and affective items** and interrogate the somatic ones instead of scoring them. Anhedonia in activities that make no physical demand is the most useful single question on a cardiology ward: a person too breathless to walk to the gate can still say whether the radio or a visitor gives any pleasure. Guilt, worthlessness, hopelessness out of proportion to the prognosis actually given, and the thought that others would be better off without them are produced by mood.

Timing is the strongest discriminator on history. If breathlessness and oedema improved with cardiac treatment while mood, interest and outlook stayed flat, the mood state has separated from the haemodynamic state and is behaving as an illness in its own right. If low mood arrived with a decompensation and lifted as it settled, an adjustment reaction reads better. Nocturnal waking divides the same way: waking breathless and needing to sit upright points at the circulation, while waking in the early hours, alert, with the day at its worst on rising, points at mood.

In adults with confirmed heart failure, **people with heart failure and depressive symptoms had poor health-related quality of life**, so the depressive syndrome tracks with how the person actually lives rather than with a score alone. That finding carries a direction and a population and no magnitude, so its use is to justify looking, and to justify treating what is found.

RULE **Attribute by timing and by course, never by the symptom itself.** Any single somatic symptom in heart failure is uninformative. What discriminates is whether it moved with the cardiac state or independently of it, and only a sequential history answers that. Put the cardiac events, the treatment changes and the mood changes on one timeline and the attribution usually declares itself.

8.2 Stage of illness as the organising question

Heart failure is not one psychiatric situation, and the commonest error in a viva is to describe it as though it were. The task, the assessment priority and the realistic intervention all change with the stage.

STAGE	THE PSYCHOLOGICAL TASK	WHAT THE ASSESSMENT MUST REACH
New diagnosis	Absorbing a prognosis and a changed identity, often while still physically well	What they were told, what they believe it means, and who else heard it
Stable on treatment	A daily regimen, fluid and salt restriction, weighing, symptom diaries	Adherence, the cost of the regimen, and whether the effort has become hopeless rather than tiresome
Recurrent decompensation	Repeated admission, unpredictability, lost confidence in the body	Demoralisation, fear of dying in an episode, sleep, the carer's own state
Advanced illness	Anticipated dying, dependence, the meaning of continued treatment	Depression as distinct from a considered response to a real prognosis

The hand-offs are explicit. Demoralisation as a construct, and the screening and treatment of depression and adjustment in a life-limiting diagnosis, are taught in the organ-interface routing section, which also owns palliative care psychiatry and the desire for hastened death. Delirium in the decompensated or dying cardiac patient belongs to the stroke, traumatic brain injury and focal brain syndromes notes, including the bedside instruments.

Cognition is examined rarely and matters constantly. Where attention and processing speed are slowed in advanced disease, a crowded regimen and a consent conversation are at risk, and the risk peaks during a decompensation. Since **heart failure with depressive symptoms ran with poor health-related quality of life**, the stage review is a functional review as much as a diagnostic one.

8.3 Arrhythmia in the psychiatric clinic, and the separation from panic

Arrhythmia reaches psychiatry along several routes: a paroxysmal rhythm disturbance already called panic, an anxiety syndrome that grew out of a documented arrhythmia, a rhythm disturbance and an anxiety disorder present together, and the causes shared by each, which include alcohol, stimulants, thyroid disease and sleep deprivation. The separation is made on history.

Onset and offset carry most of the information. A paroxysmal rhythm disturbance tends to start and stop abruptly, sometimes with a trigger such as a change of posture, and the racing heart precedes the fear. A panic attack builds to a peak over minutes and settles more gradually, fear and bodily symptoms rise together, and catastrophic thinking about the heart is present at the peak. Episodes that wake a person from sleep, that carry syncope, or that occur on exertion rather than at rest shift the balance towards a cardiac cause.

TRAP **Diagnosing panic disorder in a patient whose rhythm has never been recorded during an episode.** They overlap almost completely at the level of reported symptoms, and the label is sticky: once panic is written, every later episode is read through it and the rhythm is never captured. The requirement is an ambulatory or event recording during a typical attack, and the formulation should state whether one exists.

The reverse error is commoner in cardiology. A person with a documented rhythm disturbance and an anxiety disorder is often told the palpitations are anxiety, which is heard as an accusation and produces exactly the hypervigilant self-monitoring that amplifies the next episode. The workable formulation is that the rhythm is real, the fear is real, and each raises the salience of the other.

8.4 Rhythm procedures and the peri-procedural window

Atrial fibrillation is an episodic, bodily felt and unpredictable illness, and unpredictability is the property that generates anticipatory anxiety in any organ system. The load in the group referred for rhythm procedures is substantial: **depressive and anxiety disorders are highly prevalent among patients with atrial fibrillation (AF) undergoing pulmonary vein isolation (PVI)**.

Because **depressive and anxiety disorders are highly prevalent in the group referred for pulmonary vein isolation**, screening belongs before the procedure rather than after a disappointing result. A person anxious before ablation who remains symptomatic after it is easily written off as a treatment failure when the rhythm record is clean. Naming it in advance changes how a residual symptom is read: as an expected part of recovery rather than as proof the procedure failed.

Symptom-rhythm concordance is the idea underneath this. Do the symptoms the patient reports coincide with the rhythm the record shows? Concordance supports a rhythm intervention; discordance says something other than the rhythm is generating the experience, usually interoceptive vigilance, which responds to psychological work rather than to further ablation.

8.5 Living with an implanted defibrillator

An implanted defibrillator is unlike any other treatment a psychiatrist meets, because its therapeutic action is an aversive event delivered without warning from inside the body. The psychopathology follows from that: the therapy hurts, its timing is unpredictable, and it cannot be escaped by leaving the situation. A sudden painful event paired with whatever the person happened to be doing is all ordinary conditioning needs.

Shock-related anxiety is anticipatory: continuous apprehension about the next discharge, with hypervigilant monitoring of the pulse, the chest and any sensation that preceded a previous one. **Avoidance** is the disabling part and the part most often missed, because the person rarely volunteers it. Exertion, sexual activity and driving are avoided in case they provoke the device, being alone in case a discharge occurs with nobody present, and public places because a discharge witnessed by strangers humiliates as well as frightens. After repeated discharges in a short period, intrusive recollection, hyperarousal and phobic avoidance can take a post-traumatic form. Some

patients describe a **phantom shock**, the vivid experience of a discharge the device record shows did not happen.

The household is part of the presentation. Partners commonly witness the discharge, and witnessing what looks like a resuscitation produces its own distress, its own hypervigilance and its own protective restriction of the patient's activity, so a couple can arrive at a shared avoidance that neither of them has named. Hence the recommendation that **psychological assessment, monitoring, and therapy should be offered to ICD patients and their partners as part of routine care**. Each noun does work: assessment at implantation and after any discharge, monitoring because a person well at the device clinic one month may be housebound the next, therapy because detection alone is not a service, and partners because the couple is the unit of care.

RULE **Ask what the person has stopped doing since the device was implanted.** Anxiety is reported readily and avoidance is volunteered almost never, yet avoidance is what removes exercise capacity, work, driving, intimacy and social life. Put the question to the patient and to the partner separately.

TRAP **Reading shock anxiety as simple reassurance-seeking and answering it with reassurance.** Reassurance relieves for minutes and teaches the person that the fear requires an external answer, so the checking increases. The device is genuinely capable of firing, so honest information plus a rehearsed plan for what to do when it fires converts an unbounded threat into an event with a script.

8.6 The device at the end of life, and the deactivation conversation

As cardiac disease advances, the shock function can fire during the dying process, with no prospect of restoring a life the person is already leaving. Deactivating the shock therapy is therefore a planned conversation, held while the person is well enough to hold it and recorded where the treating teams can find it. Psychiatry is asked in for three reasons: capacity, whether depression is colouring the decision, and how to hold the discussion with a family who may hear deactivation as abandonment.

Capacity here is decision-specific and time-specific, so the question is whether this person can make this decision now, and the answer can change with the delirium, the hypoxia or the sedation of an acute episode. The doctrine and the statutory test belong to the mental health legislation notes. What psychiatry adds is the mood assessment underneath: a wish to stop a burdensome therapy late in an illness can be a considered choice, whereas the same wish carried on pervasive hopelessness and a conviction of being a burden requires that the depression be treated and the question asked again.

Families frequently hear deactivation as switching the device off altogether, or as an act that causes the death. What deactivation covers is for the cardiology and device team to answer explicitly, in the family's own language; the psychiatric contribution is to ensure the question is asked and the answer understood.

8.7 Non-drug management, and who does what

Everything here is deliverable without a prescription, which matters where the cardiac regimen is already crowded.

- A written plan for what to do if a discharge occurs or an episode starts: rest, position, who to telephone, when to attend, and when the device record will be checked.
- A graded return to activity, agreed with the cardiologist in writing and shared with the partner. Avoidance is reversed by activity that has been sanctioned; encouragement without sanction moves the conflict into the household.
- Cognitive-behavioural work aimed at the maintaining cycle: catastrophic interpretation of a bodily sensation, checking, avoidance, and the relief that reinforces it. Interoceptive exposure needs explicit cardiology agreement first.
- Paced breathing and applied relaxation, taught as skills for the episode rather than offered as a cure.
- Couple-inclusive sessions as standard rather than as an extra, since **psychological assessment, monitoring and therapy are recommended for defibrillator patients and their partners as part of routine care**.
- Coordination with cardiac rehabilitation where it exists, and with the device clinic in every case, including a question about remote monitoring: it reassures some patients and feeds the vigilance of others.

Sleep is where the cardiac and psychiatric arms of the plan meet. Orthopnoea, nocturia from the regimen, fear of dying in sleep and anticipatory fear of a night-time discharge each fragment the night and each has a different remedy, and asking what wakes the person, and what they do next, separates them faster than any sleep questionnaire.

What belongs elsewhere, named: antidepressant choice, prognostic effect and cardiac safety sit with depression and anxiety after myocardial infarction and in coronary artery disease; organ-adjusted dosing with the psychopharmacology notes; the acutely stress-precipitated cardiomyopathy that mimics infarction with the stress-heart interface section; delirium with the stroke, traumatic brain injury and focal brain syndromes notes; palliative care psychiatry with the organ-interface routing section.

Timing, not the symptom WHAT DECIDES ATTRIBUTION	**Stage, not the label** WHAT ORGANISES THE ASSESSMENT	**Recorded, not assumed** WHAT THE PANIC LABEL NEEDS FIRST	**Avoidance, and the partner** WHAT THE DEVICE INTERVIEW MUST REACH

Each tile carries a decision rather than a fact, and not one of them can be settled from the symptom in front of you: attribution comes from the timeline, priority from the stage, the panic label from a recording, and the device interview from the partner as well as the patient.

IN INDIA

Device therapy is largely an out-of-pocket purchase, the device clinic may be a day's travel away, and the family rather than the individual usually carries the decision and the follow-up. The couple-inclusive element of the recommendation is therefore easier to deliver here than in services built around the individual patient, while longitudinal monitoring is harder, because the person is seen when they can travel rather than when they are due. A written plan the family can act on at home does more work here than a scheduled review. What this section does not hold is worth stating: no Indian prevalence figure for depression in heart failure, no Indian count of implanted devices, no Indian cost figure for defibrillator therapy.

GOOD TO REMEMBER

A short set of questions does most of the ward work: what has changed since the last cardiac event and in what order, what the person has stopped doing since then, and who else was in the room when the prognosis or the device was explained. Attribution, avoidance, and the partner who is carrying the other part of the distress.

The plate lays the limbs of the assessment side by side: the shared somatic symptoms with the questions that separate them, the arrhythmia presentations and where each routes, the psychopathology of living with a device, and the non-drug management that follows. Prescribing is absent from it by design.

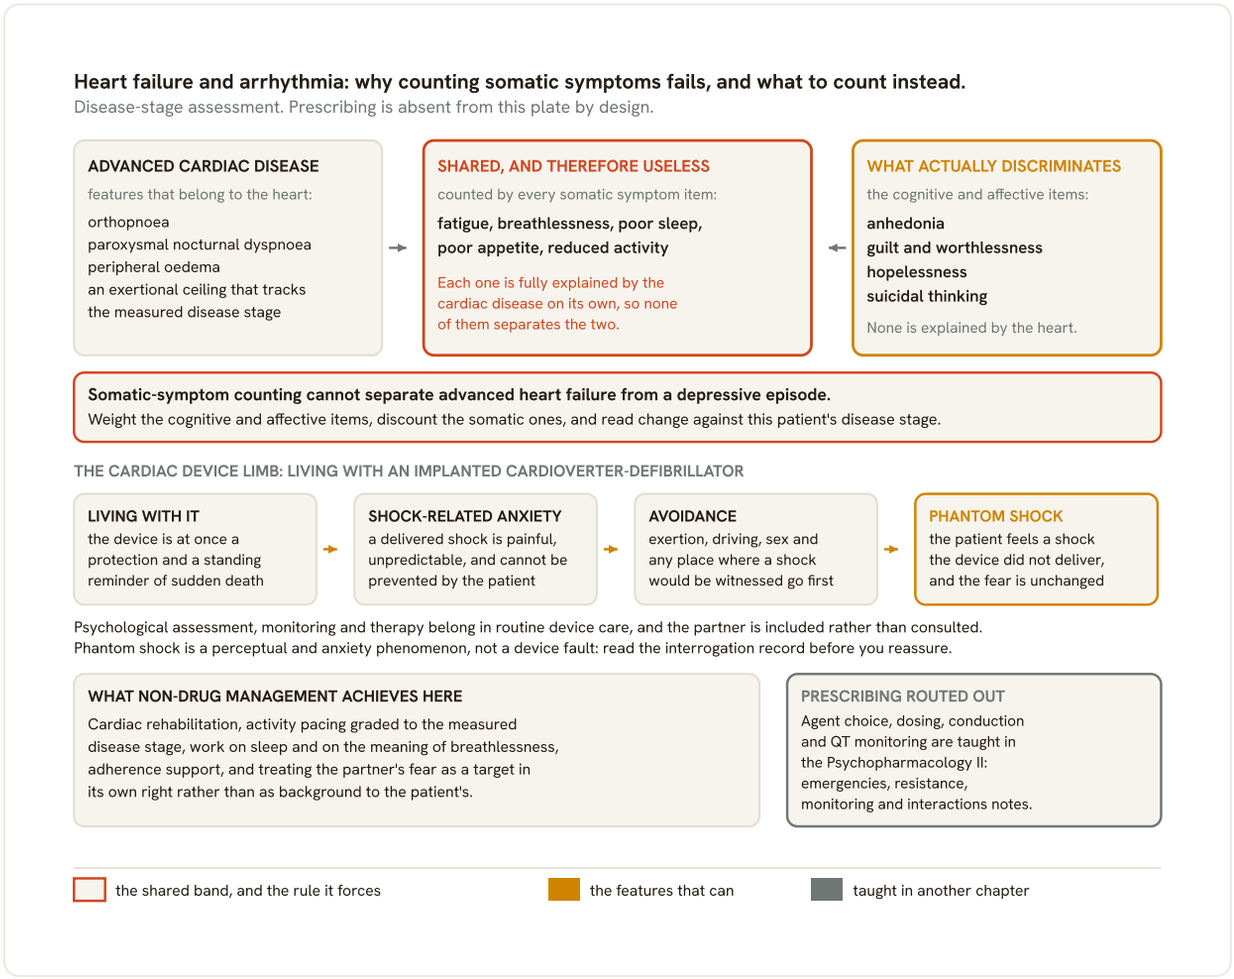


Fig 40.2 Disease-stage assessment in heart failure and arrhythmia, built around the symptom-overlap problem. The centre panel holds the five features a somatic symptom count would score, and every one of them is fully produced by advanced cardiac disease on its own, which is why an instrument weighted toward somatic items inflates in this population and a candidate who counts them reaches the wrong answer confidently. The discrimination lives in the right-hand panel, in items the heart does not generate. The device limb runs on its own logic: the implanted cardioverter-defibrillator protects and threatens at the same time, a delivered shock is painful and cannot be prevented by the patient, avoidance follows, and phantom shock closes the loop by producing the fear without the event. Prescribing is deliberately absent. From memory you should be able to name the five shared features, the four discriminating ones, and the four steps of the device limb, and say why the first list cannot settle the diagnosis.

> Attribute by timing and course, ask what the person has stopped doing since the device was implanted, and treat the couple, since psychological assessment, monitoring and therapy should be offered to defibrillator patients and their partners as part of routine care.

9 · Psychiatric aspects of COPD, asthma and chronic respiratory failure

Breathlessness and panic recruit the same body, and the patient living through either one cannot reliably tell you which of them started. Respiratory psychiatry is examined on that ambiguity and on what a clinician does about it at the bedside. This section installs four things: a working discrimination between the dyspnoea of obstructive airways disease and the dyspnoea of a panic attack; the two specialty-guideline positions this chapter can anchor to a respiratory source; the behavioural levers that move outcome, which are referral to rehabilitation and inhaler adherence; and an honest declaration about chronic respiratory failure, where a specialty

statement was sought and not obtained. Cardiac material is not taught here. Heart failure, arrhythmia and cardiac devices have their own section, and the psychiatric sequelae of critical illness, including post-intensive-care syndrome, sit in the merged organ-interface routing section.

9.1 The discrimination this section exists for

A patient who is breathless, frightened, tachycardic, tremulous and unable to complete a sentence has produced a presentation that two very different processes can generate. Airflow limitation produces breathlessness, and breathlessness produces fear; fear produces hyperventilation, and hyperventilation produces breathlessness. The loop closes in both directions, which is why the question is rarely which one the patient has, and almost always which one is driving the episode in front of you.

Several features shift the probability without settling it. The temporal shape is the most useful single discriminator: a panic attack crescendos and then subsides over a bounded period, whereas an exacerbation of airways disease climbs and stays. The relation to exertion matters, because bronchoconstriction and deconditioning both load onto effort while a panic attack characteristically arrives at rest and out of context. Objective airflow limitation, audible wheeze, a fall in oxygen saturation and a response to a reliever inhaler point one way; circumoral and peripheral paraesthesiae, carpopedal spasm and the light-headedness of hypocapnia point the other, since an obstructed patient usually cannot blow off enough carbon dioxide to produce them. The content of the thought is informative and is often the most accessible part of the history: the catastrophic cognition in panic is about dying or losing control now, and it recurs in the same form.

None of this licenses a clean split, and the examinable position is that it does not need to. A patient with obstructive lung disease who develops recurrent panic attacks has both conditions, each worsening the other, and treating only one leaves the disability where it was. **Anxiety sensitivity**, the tendency to read bodily sensations as dangerous, is the variable that couples them, and it is the one a psychiatrist can treat.

RULE **The dangerous direction of error is treating airflow limitation as anxiety, not the reverse.** Calling a panic attack an exacerbation costs an unnecessary course of treatment. Calling an exacerbation a panic attack can cost a life. Where the two cannot be separated in the room, treat the airway first, interrogate the psychiatry afterwards, and record why.

9.2 Why mood and anxiety cluster in obstructive lung disease, and what they cost

The association is not a coincidence of two common conditions sharing a population, and a candidate who says only that will be pressed. Several distinct pathways converge, and naming them separately earns the marks.

The behavioural pathway is the strongest and the most modifiable. Breathlessness on exertion is aversive, so the patient reduces exertion; reduced exertion produces deconditioning, and deconditioning lowers the threshold at which the next effort produces breathlessness. The spiral narrows the patient's world by degrees and delivers the ingredients of depression, which are loss

of activity, loss of role and loss of contact. From the anxiety side it is avoidance learning: exertion is paired with a frightening sensation, and the patient stops.

Beyond that sit contributions this chapter can name but does not quantify. Tobacco use is a shared risk factor for the lung disease and for several psychiatric conditions, and it means a proportion of these patients carry a nicotine dependence any plan must address. Chronic hypoxaemia and hypercapnia affect attention, processing speed and daytime alertness, and disrupted nocturnal breathing fragments sleep, which independently degrades mood. Systemic corticosteroid exposure is common here and carries its own psychiatric consequences; that construct is owned elsewhere and is named only as a hand-off. Finally, the disease is progressive, visible and frightening with a foreseeable trajectory, and demoralisation in that setting is a reasonable response as well as a clinical problem.

The cost is not only distress. Comorbid mood and anxiety symptoms track with poorer adherence, more frequent presentation to acute services and worse functional status at any given level of measured airflow limitation. No prevalence figure, effect estimate or readmission proportion is held by any claim row available here, so none is stated: carry the direction, and say openly that the magnitude was not obtained.

TRAP **Deciding that low mood in advanced lung disease is understandable, and therefore leaving it alone.** Comprehensibility is not a diagnostic exclusion, and this reasoning would not be accepted for any other comorbidity. It matters more than usual here because the depressive component is what removes the patient from the activity, the rehabilitation and the adherence that would change the trajectory. An understandable depression and a treatable depression are the same depression.

9.3 Pulmonary rehabilitation, and the referral statement to memorise

Pulmonary rehabilitation is a structured programme of supervised exercise training combined with education and self-management support for patients with chronic respiratory disease. It belongs in a psychiatry chapter because it acts on the behavioural spiral above: it is the intervention most likely to reverse the deconditioning, avoidance and social withdrawal carrying much of the psychiatric burden here.

The single most examinable statement this section owns concerns who gets referred. The specialty guideline position is that **coexistent symptoms of anxiety and/or depression in patients with COPD should not preclude referral to pulmonary rehabilitation**. That sentence is worth learning in its shape as well as its content, because it is written as a prohibition on exclusion rather than as an encouragement to refer.

It is written that way because the pressure in real services runs the other way. A patient who appears low, unmotivated or unlikely to complete a course is exactly the patient a stretched programme is tempted to screen out, and the screening is usually informal, done by the referrer, on a judgement about probable engagement rather than by any stated criterion. The guideline blocks that move, and the logic is worth stating in a viva: the psychiatric symptoms are in

substantial part a product of the very deactivation the programme treats, so using them to withhold it removes the intervention from the group with most to gain.

Two boundaries keep the statement honest. It concerns referral and not completion, so supporting attendance, which is where anxious and depressed patients genuinely fall away, remains a clinical task. And no effect size, completion rate or symptom change accompanies the recommendation in the row this chapter holds, so no magnitude is claimed. What a candidate can say without qualification is that coexistent anxiety and depressive symptoms are not a reason to withhold referral to pulmonary rehabilitation, and that the psychiatrist's contribution is often to say so to the referring team.

RULE **Psychiatric comorbidity is an indication to support the referral, not a filter applied before it.** When a respiratory team asks whether a low or anxious patient is suitable for rehabilitation, the guideline-anchored answer is that coexistent symptoms of anxiety and/or depression should not preclude referral. The useful follow-up question is not whether to refer, but what must be arranged so the patient reaches the second session.

9.4 Adherence and technique, which is the behavioural half of disease control

Inhaled treatment is the one therapeutic area where a psychiatrist's ordinary skills change a respiratory outcome, because almost every failure of inhaled treatment is behavioural rather than pharmacological. Two distinctions organise the assessment. The first is between **unintentional non-adherence**, where the patient intends to take the treatment and does not manage it, and **intentional non-adherence**, where the patient has decided against it. The second is between not taking the drug and not delivering it, since a patient with poor inhaler technique is fully adherent and receiving very little.

Psychiatric symptoms map onto these categories specifically rather than generally. Depressive symptoms produce forgetting, loss of routine and the hopeless reasoning that treatment will not help anyway. Anxiety produces overuse of the reliever inhaler and a dependence on it that masks deteriorating control and intermittently reinforces using it again. Beliefs about harm, particularly the fear that inhaled steroid preparations are dangerous or habit-forming, produce quiet intentional non-adherence the patient will not volunteer. Cognitive impairment and manual difficulty produce technique failure independent of all of this.

FAILURE MODE	WHAT IT LOOKS LIKE IN CLINIC	WHAT ACTUALLY CHANGES IT
Technique failure	Correct prescription, exhausted canister, poor control, patient reports full adherence	Watching the patient use the device, then re-teaching or changing the device
Forgetting and lost routine	Erratic use, worse in the morning, tracks with mood and with disorganised days	Anchoring the dose to an existing fixed daily act, plus treating the mood state
Reliever overuse	Frequent reliever use, escalating anxiety, preventer untouched	Naming the cycle explicitly, then treating the anxiety alongside the airway
Belief-driven refusal	Articulate reasons, often about steroids, offered only when asked directly	Eliciting the belief without correcting it first, then addressing the specific fear

The instruction that follows is single and easy to examine: ask the patient to demonstrate, do not ask the patient to report. A demonstration takes less time than the conversation about escalating treatment that follows from assuming the technique was adequate.

Dyspnoea against panic: six discriminators, and the band where they genuinely overlap.
Neither outer column settles the question alone. The middle column is the reason it stays hard.

FEATURE	DYSPNOEA OF COPD OR ASTHMA	WHERE THEY OVERLAP	PANIC ATTACK
ONSET	over minutes to hours, from a baseline that was already limited	both begin with breathlessness and fear, in the same body	abrupt, peaking within minutes from an unremarkable baseline
TRIGGER	exertion, cold air, allergen, infection or a missed preventer	exertion can start either, and fear can follow either	often at rest and out of context, or cued by fear of an attack
TIME COURSE	climbs, and stays up until the airway is treated	a long attack and a short exacerbation look alike	crescendos, then subsides over a bounded period
WHAT RELIEVES IT	a reliever inhaler, and treating the airway itself	reliever overuse eases both for a while, and masks both	slowed breathing and time: the attack is self-limiting
WHAT THE PATIENT DOES	drops exertion, then drops the activities that needed it	avoidance narrows the life the same way from either cause	escapes the situation, carries the inhaler as a safety object
WHAT THE OBSERVER SEES	wheeze, accessory muscle use, falling oxygen saturation	tachypnoea, tachycardia, tremor and visible fear in both	circumoral and peripheral paraesthesiae, carpopedal spasm

The dangerous direction of error is treating airflow limitation as anxiety, never the reverse.
Where the two cannot be separated in the room, treat the airway first, interrogate the psychiatry after, and record why.

LEVER ONE: PULMONARY REHABILITATION
Coexistent symptoms of anxiety and/or depression in patients with COPD should not preclude referral to pulmonary rehabilitation. It is written as a prohibition on exclusion, not an encouragement.

LEVER TWO: INHALER ADHERENCE AND TECHNIQUE
Separate not taking the drug from not delivering it, and unintentional from intentional non-adherence. Then ask the patient to demonstrate the device rather than to report on it.

ROUTED OUT Steroid-induced psychiatric disturbance is taught in the Endocrine and metabolic psychiatry notes.
ROUTED OUT Agent choice, dosing and monitoring are taught in the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes. This plate names no agent and prints no dose.

Fig 40.3 Separating the dyspnoea of obstructive airways disease from a panic attack, with the overlap drawn rather than glossed. The two outer columns are the features that shift probability, read across six axes rather than argued one at a time. The coral column is the honest part: breathlessness and fear recruit the same body, the loop closes in both directions, and a patient with obstructive lung disease who also has panic attacks has both conditions, each worsening the other. That is why the examinable question is almost never which one the patient has but which is driving this episode and in what proportion. The two levers below are where a psychiatrist changes a respiratory outcome, and both are behavioural rather than pharmacological. From memory you should be able to run the six axes in either direction, state the asymmetry of risk in the rule box, and give the referral sentence and the demonstration instruction verbatim.

9.5 Asthma in adolescence, and the association that has to be stated carefully

Asthma carries a second and separate examinable position, and it is age-specific. The specialty guideline records that **asthma in adolescence is associated with an increased likelihood of major depression, panic attacks and anxiety disorder**. Three things about that sentence are worth holding. It names adolescence rather than asthma at any age. It names three outcomes rather than a general elevation in distress. And it states an association, so a candidate who converts it into a causal claim in either direction has overstated the source.

The direction genuinely is uncertain and saying so is the stronger answer. A chronic, visible, sometimes frightening condition that constrains sport, sleep and social life at the stage where all three matter most is a plausible cause of low mood. Anxiety, equally, worsens the perception of breathlessness and the behaviours that control it. Shared factors sit underneath both, including disrupted sleep, corticosteroid exposure and family stress.

The consequence for practice is concrete. Because adolescent asthma is associated with an increased likelihood of major depression, panic attacks and anxiety disorder, mood and panic symptoms are worth asking about routinely in this group rather than waiting for them to be raised, and mood and anxiety are candidate contributors to poor asthma control rather than a separate matter for later. The overlap with panic is awkward because both under-perception and over-perception of airflow limitation occur: the adolescent reporting severe symptoms may be obstructed, panicking, or both, and the one reporting little may be seriously obstructed.

TRAP **Labelling an adolescent's attack as panic and thereby delaying respiratory treatment.** The association above makes panic more likely in this group, which is exactly why the label becomes available and exactly why it is dangerous. An adolescent with asthma presenting with an acute frightening episode is treated for the airway on clinical grounds, and the psychiatric formulation is completed once the episode has settled. Offering reassurance as the first move in an acute presentation answers the question badly, however well the comorbidity is discussed afterwards.

9.6 Chronic respiratory failure, and the gap this chapter declares

The third condition in this section's title is the one for which no guideline statement was obtained, and the honest position has to be stated rather than papered over. A specialty-guideline statement specifically linking chronic respiratory failure with psychiatric sequelae, with psychiatric assessment, or with psychiatric management was sought for this chapter from the respiratory bodies whose statements would carry such a recommendation, being the national thoracic guideline body and the international obstructive lung disease initiative. It was not obtained in a form this chapter can bind. That absence describes what this chapter opened. It is not evidence that no such statement exists, and it must not be converted into a claim that the specialty is silent.

What a candidate does without it is reason from general principles and say that this is what they are doing. Chronic hypercapnia produces daytime somnolence, headache, impaired attention and, as it advances, confusion, any of which can present first as an apparent psychiatric change. Nocturnal hypoventilation fragments sleep, which degrades mood independently. Long-term oxygen therapy and domiciliary ventilation impose visible dependence on a machine and reorganise family life around the equipment. The loss of control is the psychological core, and interventions that return any control to the patient tend to help.

Two boundaries apply. Acute confusional states here are assessed and managed under the delirium material in the stroke, traumatic brain injury and focal brain syndromes notes, and no bedside delirium instrument, rating scale or antipsychotic choice is taught in this section. And a patient declining domiciliary ventilation raises a capacity question; capacity is decision-specific and time-specific, and everything beyond that, including the statutory test, belongs to the mental health legislation notes.

IN INDIA

No figure for the availability, distribution or coverage of pulmonary rehabilitation programmes in India is held by any claim row available to this chapter, so none is offered. What can be said is what the referral position implies for a service that has a programme and for one that does not. Where a programme exists, coexistent anxiety and depressive symptoms are not grounds to withhold referral to pulmonary rehabilitation, and a psychiatrist asked to assess a low patient before referral has a guideline-anchored answer for the respiratory team. Where no programme exists, which is the position across much of the country outside larger centres, the deliverable residue is still the active ingredient: a graded, written, supervised increase in activity that a physiotherapy service, a nursing team or an informed family member can sustain, with education about the breathlessness cycle. The gap in a district service is usually the structure and the supervision, not the exercise.

9.7 What this section hands off, and does not teach

Corticosteroid-induced psychiatric disturbance is common in this population and is owned elsewhere. It is named here as a differential whenever mood, sleep or mental state changes after a course of systemic steroid; the mania, depression, psychosis and management content sits with the corticosteroid owner in the endocrine and metabolic material and on the chapter's routing plate.

Prescribing in respiratory disease routes to the psychopharmacology notes without exception. Sedative burden and respiratory depressant risk are why a psychiatrist is cautious here, and the drug-by-drug reasoning, the comparative rankings and every dose belong there, so none is restated. The same applies to the effect of tobacco smoke on the metabolism of certain psychotropics: it is named as a hand-off, and neither its direction nor its magnitude is stated here.

Non-pharmacological treatment of the anxiety component follows ordinary principles with one adaptation. Interventions that deliberately provoke the feared bodily sensation require care when the airways are genuinely obstructed, because the sensation is not false. The workable version treats the interpretation and the avoidance rather than trying to show the sensation is harmless, and it runs alongside optimised inhaled treatment.

GOOD TO REMEMBER

On the ward, three moves cover most of the value a psychiatrist adds to a breathless patient. Watch the inhaler being used before believing anything about adherence. Ask what the patient thinks is happening during an attack, because the answer separates the catastrophic cognition of panic from the accurate fear of someone obstructed before. And ask whether the patient has been referred for rehabilitation, because anxiety and depressive symptoms are not a reason to preclude that referral and a psychiatric opinion is often what is used to justify not making it.

9.8 The answer, assembled

An examiner asking about psychiatric aspects of respiratory disease wants a structure rather than a list. The following sequence carries this section.

- State the discrimination first: breathlessness and panic share a physiology, the loop runs both ways, and the common answer in this population is both rather than either.
- Give the safe rule of error: treat the airway first when the two cannot be separated, because the costly mistake is calling obstruction anxiety.
- Explain the behavioural spiral of breathlessness, avoidance and deconditioning, since it accounts for much of the presentation and is what treatment targets.
- Deliver the referral position verbatim: coexistent symptoms of anxiety and/or depression in patients with COPD should not preclude referral to pulmonary rehabilitation.
- Deliver the asthma position with its age qualifier and its association wording: asthma in adolescence is associated with an increased likelihood of major depression, panic attacks and anxiety disorder, and say the direction was not established.
- Cover adherence and technique as a behavioural target, and demonstrate rather than ask.
- Declare the chronic respiratory failure gap openly: reason from hypercapnia, sleep, dependence and loss of control, and say that a specialty statement binding psychiatric management to chronic respiratory failure was sought and not obtained.

Both, not either	**Refer, not exclude**	**Watched, not reported**	**Declared, not inferred**
WHAT DYSPNOEA PLUS PANIC USUALLY MEANS	WHAT COEXISTENT LOW MOOD CHANGES	HOW INHALER ADHERENCE IS ASSESSED	WHAT TO SAY ON RESPIRATORY FAILURE

Each tile names a default that ordinary practice sets the other way, which makes the strip an audit of one's own clinic behaviour rather than a summary of the text.

PITFALL

Answering this item as a general essay on the psychological impact of chronic illness. The material is true and it earns very little, because it could have been written about any organ. What makes an answer respiratory is the dyspnoea and panic discrimination, the referral position on rehabilitation, the age-specific asthma association and the behavioural analysis of inhaled treatment.

Breathlessness and panic are usually both present; treat the airway first, refer for rehabilitation regardless of coexistent anxiety or depression, watch the inhaler being used, and ask every adolescent with asthma about mood and panic.

10 · Takotsubo cardiomyopathy and the stress-heart interface

Takotsubo syndrome is the plainest demonstration in medicine that an acute psychological event can produce a structural cardiac lesion, which is why it belongs in a psychiatry paper and not only in a cardiology one. An emotional or physical stressor is followed, usually within hours, by chest pain, breathlessness or collapse, by an electrocardiogram reading as an evolving anterior infarct, by rising biomarkers of myocyte injury, and by a left ventricle gone akinetic over

a territory larger than any single coronary artery supplies. The arteries are then found unobstructed, and the wall-motion abnormality resolves, which infarcted muscle never does. This section teaches the syndrome by mechanism rather than by table: the epidemiological specifics a resident would ordinarily quote are absent from this build's rows, and are declared rather than supplied from memory.

Several neighbouring topics are named here and taught elsewhere, and knowing the boundary is itself examinable. Depression and anxiety after myocardial infarction, and safe antidepressant choice in established coronary disease, belong to the organ-interface routing section. Heart failure, arrhythmia and implanted cardiac devices have their own section. Delirium on the coronary care unit sits with the notes that own delirium, and the psychological sequelae of critical illness with intensive-care psychiatry. Organ-adjusted dosing is a psychopharmacology question. What remains is the mechanism, the differential and the liaison contribution.

10.1 The shape of the syndrome, and why a psychiatrist is called at all

The syndrome is defined by a pattern rather than by any single test, and each element carries weight. There is an acute regional abnormality of left ventricular contraction, extending beyond the distribution of one coronary artery, which makes an occlusive explanation implausible on anatomical grounds alone. There is no culprit obstruction able to account for it. And there is recovery, so the ventricle akinetic on admission contracts normally later. Transience is not a happy outcome appended to the definition, it is part of the definition, and why the diagnosis is partly retrospective.

The name describes the shape the ventricle takes at end systole in the commonest form, a rounded akinetic apex above a narrow hypercontractile base, resembling a narrow-necked pot. That image encodes the paradox the mechanism has to explain: the segment furthest from the base fails while the base works harder than usual.

A psychiatrist is called into this picture for a small number of reasons, and knowing them in advance determines what a useful consultation note contains. The antecedent was psychological and somebody wants it characterised. A psychotropic drug, a substance or a withdrawal state is suspected. The patient is distressed or agitated on a cardiac unit. Or the team, having found clean arteries, is uncertain what to tell the patient and what to arrange on discharge. Each is answerable. Reciting a cardiology description back to a cardiologist is not.

10.2 The catecholamine-surge account, taken in the order the account itself runs

The dominant mechanistic account is a **catecholamine surge** acting on myocardium, and it is best learned as a chain, because each link is separately testable and separately arguable.

The chain begins centrally. A stressor is appraised as threatening, or the body is placed under acute physiological load, and the central autonomic network drives the sympatho-adrenal axis. The adrenal medulla releases adrenaline into the circulation; cardiac sympathetic terminals release noradrenaline into the myocardial interstitium. These are separate deliveries with different distributions, which matters later: circulating adrenaline reaches the whole ventricle, whereas neural noradrenaline concentrates wherever innervation is densest.

The next link is the effect of that load on the myocyte. High catecholamine exposure drives calcium entry and calcium overload, and calcium-overloaded myocytes contract in a disordered, sustained way. The histological correlate described in this setting is **contraction band necrosis**, a pattern distinct from the coagulative necrosis of coronary occlusion, accompanied by oxidative stress, mitochondrial dysfunction and falling cellular energy supply. This is a toxic injury, not an ischaemic one, and the distinction is why the muscle can recover.

A parallel link is vascular. Catecholamine excess is associated with epicardial coronary spasm and, more importantly here, with **microvascular dysfunction**: impaired dilatation of the small vessels that determine actual myocardial perfusion. Perfusion can therefore fall while the large arteries remain widely patent, which reconciles a genuinely injured myocardium with a normal angiogram.

The most interesting link, and the one that best explains why the apex, is receptor level. Adrenoceptor density is described as higher towards the apex while sympathetic innervation is sparser there, leaving the apex disproportionately exposed to circulating adrenaline. The subtype predominating apically can also switch its intracellular coupling: at ordinary concentrations it signals through the stimulatory G protein and increases contraction, and at the concentrations reached in a surge it couples instead to the inhibitory G protein, reducing contraction. On this account the apex is not merely overwhelmed, it is actively switched off, and the switch is protective, sparing the myocyte a lethal calcium load at the cost of regional akinesia.

The final link is **myocardial stunning**. Stunned myocardium is alive but not contracting: the cells persist, the contractile apparatus is temporarily disabled, and function returns when the insult is withdrawn. Stunning is why a ventricle that looks infarcted on admission can look normal later, and it is the single concept that makes the whole syndrome coherent.

RULE **Learn the chain, not the label.** Central appraisal, sympatho-adrenal outflow, adrenal and neural catecholamine delivery, calcium overload with contraction band necrosis, microvascular impairment, apical receptor coupling switch, regional stunning, then recovery. A candidate who can walk that chain can answer any question asked about this syndrome, including those on variants and differential diagnosis, because each attacks a specific link.

10.3 What the account explains well, and where it strains

Good teaching of a mechanism includes its limits, and this one has clear limits. The apical receptor argument explains the commonest morphology elegantly. It explains it too well, because the syndrome also occurs with the akinetic segment midventricular, basal with a hypercontractile apex, and focal. A gradient anchored at the apex predicts apical disease; it predicts basal disease poorly. The honest reading is that the receptor gradient is one contributor within a more general catecholamine-mediated regional vulnerability, and that regional variation reflects unresolved individual differences in innervation, receptor distribution and microvascular reactivity.

The account also under-determines who develops the syndrome. Acute catecholamine surges are common in human life and this presentation is not, so a susceptibility term is required: something about vascular reactivity, hormonal environment, autonomic set point or genetic

variation separating the person who stuns from the far larger number who do not. Naming that term, and saying it is unresolved, is a better answer than asserting a single cause.

This is where a conventional account would supply a demographic profile, since the demographic skew of this syndrome is the specific most often reproduced about it. That figure is exactly what this build cannot supply: the rows available here carry no distribution by sex, no age distribution, and no proportion of acute coronary presentations accounted for by the syndrome. What can be taught instead is what such a figure would require: a consecutive unselected series taken for angiography with suspected acute coronary syndrome, an explicit diagnostic definition applied prospectively, and a stated catchment and period, so that numerator and denominator each mean something. Any figure quoted without those is a number without a population.

TRAP **Treating the receptor gradient as the whole explanation.** It is an attractive and memorable argument, and it accounts for the classical morphology only. Candidates who commit to it are then unable to explain the basal and midventricular variants, and an examiner who asks about variants is usually asking precisely this question.

10.4 Why the presentation is mistaken for an acute coronary syndrome

It is mistaken for an acute coronary syndrome because at first contact it is genuinely indistinguishable from one, which is a property of the syndrome rather than a failure of the clinician. The symptom is the same: chest pain, breathlessness, syncope or shock. The electrocardiogram carries ischaemic change in a territory, commonly ST-segment elevation anteriorly, later T-wave inversion and prolongation of the corrected QT interval. The biomarkers rise, because myocytes really are injured. The ventricle really is akinetic. Every element that ordinarily signals an occluded artery is present, and the artery is not occluded.

What separates them is therefore investigation and time, not acumen at the bedside. Angiography shows no culprit lesion. Imaging shows a wall-motion abnormality crossing coronary territories. Follow-up imaging shows recovery. None of that is available in the first minutes, which is why the correct initial management is the acute coronary syndrome pathway, and why a psychiatrist confident at the bedside that this is stress-related is confident about something they cannot yet know.

The psychiatric traps here run in opposite directions and each is serious. The first is re-labelling: the angiogram is reported clean, the word psychological enters the notes, and the presentation is reinterpreted as anxiety, panic or somatisation. Unobstructed arteries are compatible with a severely injured ventricle, with pulmonary oedema, arrhythmia and haemodynamic collapse: a clean angiogram is not a normal heart. The second is the mirror image: a recognisable emotional antecedent leads the team to a stress explanation while an atherothrombotic occlusion is in progress. A psychological trigger has no power to exclude an occlusion.

TRAP **Reading a clean angiogram as a psychiatric diagnosis.** The angiogram excludes an obstruction; it excludes nothing else. The commonest way this syndrome is mismanaged on a psychiatric liaison list is that the cardiac finding is treated as ruled out and the patient is handed over as an anxiety presentation, at which point nobody is monitoring a ventricle that is still stunned.

10.5 The psychiatric antecedents, and what can honestly be said about them

Antecedents fall into recognisable classes, and the classification here is conceptual rather than quantitative. Emotional antecedents include bereavement, acute interpersonal conflict, fear, grave news, an assault on status or reputation, and acute panic. Physical antecedents include surgery, acute severe illness, sepsis, seizure, respiratory failure and metabolic disturbance, engaging the same sympatho-adrenal pathway with no psychological content. A proportion of cases have no identifiable antecedent of either kind, which matters, because demanding one before entertaining the diagnosis will miss those patients.

A subtler point is valence. The mechanism is driven by arousal and catecholamine load, and arousal is not the exclusive property of aversive emotion, so intensely positive events have also been described as antecedents. This follows from the chain above and is not an anomaly needing separate explanation. What proportion of antecedents is emotional against physical, and what proportion is aversive against pleasurable, are quantitative questions with no grounded answer here.

Comorbidity is a different question from causation, and here this section does hold a grounded statement. **Psychiatric disorders, depression and anxiety among them, are common in patients with this syndrome**, and that is worth stating precisely for what it is: a statement of co-occurrence in the affected population. It establishes that the psychiatrist has a population-level reason to be involved rather than an incidental one.

What it leaves open is the direction of the relationship, and an examiner rewards a candidate who says so. Depression and anxiety may be an antecedent vulnerability, plausibly through sustained autonomic and hypothalamic-pituitary-adrenal dysregulation that alters the response to a surge. They may be a consequence, since surviving a frightening cardiac event with an uncertain explanation generates fear and low mood. They may reflect a shared substrate, with autonomic dysregulation producing each. Co-occurrence is compatible with all of those, and the proportion affected, with any temporal ordering, is a specific this build does not carry.

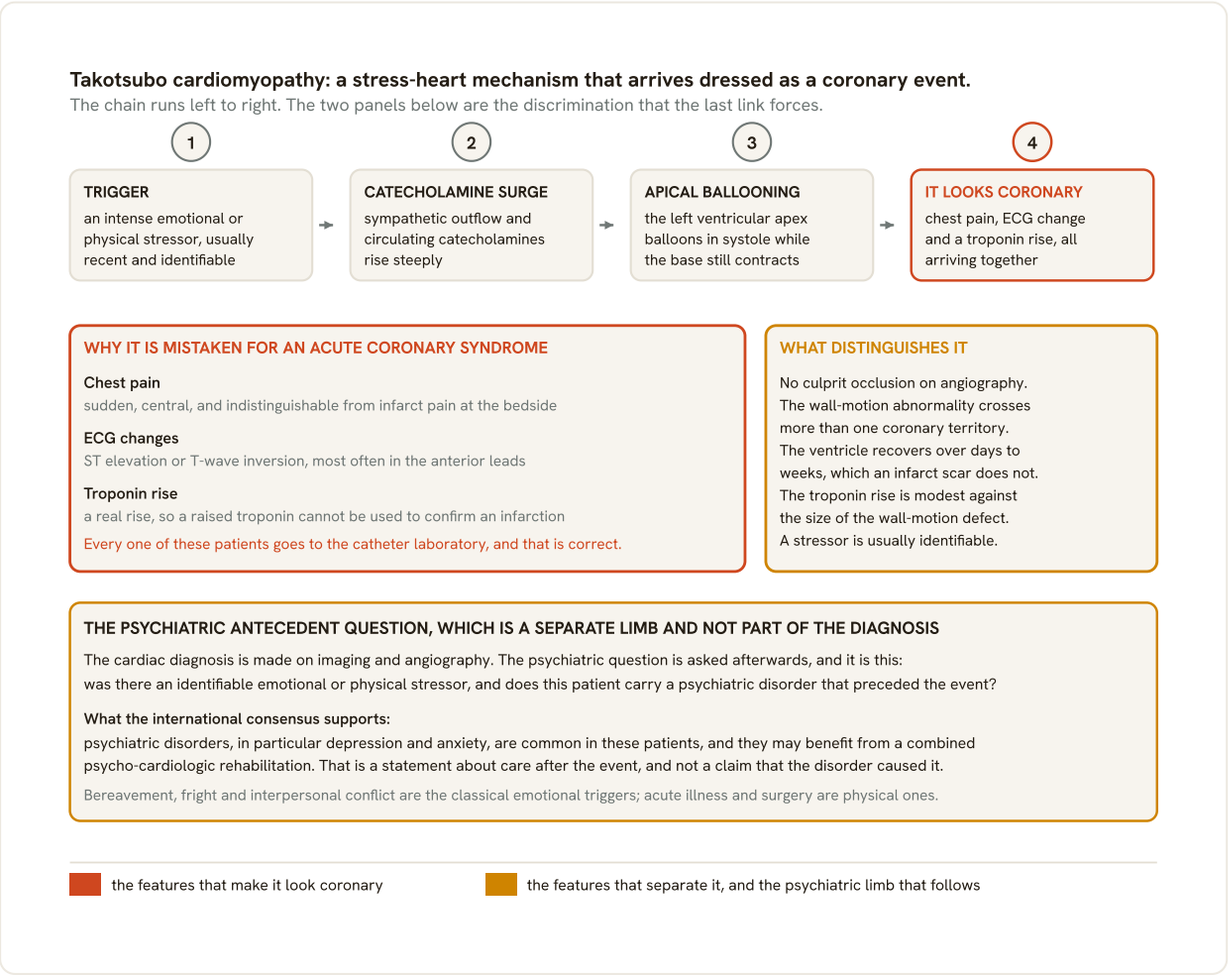


Fig 40.4 The stress-heart interface as a four-link chain, with the mimicry and the discrimination drawn side by side. The mechanism is the examinable spine: a trigger, a catecholamine surge, an apex that balloons in systole while the base still contracts, and a presentation that carries all three cardinal features of an acute coronary syndrome at once. That is why the mistake is not carelessness, and why the correct first move is still the catheter laboratory. What separates the two is the pattern rather than any single test: the wall-motion abnormality does not respect a coronary territory, no culprit occlusion is found, and the ventricle recovers in a way an infarct scar does not. The psychiatric limb is deliberately drawn apart from the diagnosis, because it is asked after the cardiac question is settled. From memory you should be able to give the four links in order, the three mimicking features, at least three discriminating ones, and the reason the antecedent question is a separate limb rather than a diagnostic criterion.

10.6 The differential of the surge, which is largely a psychiatric differential

Because the mechanism is catecholamine load, the differential of the antecedent is a list of states psychiatrists meet constantly, and this is where the liaison psychiatrist supplies information nobody else can obtain.

- **Exogenous catecholamine or sympathomimetic load.** Adrenaline given therapeutically, local anaesthetic containing adrenaline, and sympathomimetic substance use including stimulants and cocaine. The substance history establishes these.
- **Withdrawal states.** Alcohol and sedative-hypnotic withdrawal produce sustained autonomic arousal, and opioid withdrawal a lesser but real adrenergic load. A withdrawal syndrome is an endogenous catecholamine infusion the patient is unaware of receiving.
- **Iatrogenic surges in psychiatric treatment.** Electroconvulsive therapy produces a transient, deliberate and monitored autonomic surge, a recognised context in which this presentation

has been described. Management of that surge belongs to the treatment's own protocol.

- **Hyperadrenergic psychiatric emergencies.** Neuroleptic malignant syndrome, serotonin toxicity and catatonia with autonomic instability sustain adrenergic arousal for hours or days, a longer exposure than a discrete emotional trigger delivers.
- **Acute severe panic and acute stress reactions.** These are most likely to be dismissed as merely psychological on a medical ward, and the mechanism gives a reason to take them seriously as physiological events.
- **Organic causes with psychiatric presentations.** Phaeochromocytoma delivers precisely the surge the mechanism requires and may present with paroxysmal anxiety. Subarachnoid haemorrhage produces neurogenic stunned myocardium by an overlapping pathway, and prolonged seizure activity does the same. Each can reach a psychiatrist first.

Psychotropic drugs with cardiac effects, including corrected QT interval prolongation and the cardiovascular effects of stimulants, are a psychopharmacology topic, cross-referred rather than restated. Their relevance here is that the medication list is part of the surge history.

FIGURE CONVENTIONALLY QUOTED	WHAT WOULD HAVE TO BE MEASURED	WHAT THIS BUILD HOLDS
Frequency among suspected acute coronary presentations	Consecutive unselected patients angiographed for suspected acute coronary syndrome, prospective definition, stated catchment and period	Not grounded here
Sex and age distribution	Composition of a consecutive diagnosed series, read against the composition of the source population	Not grounded here
Emotional against physical antecedents	Antecedent classified prospectively by stated criteria in every case, including cases with none identified	Not grounded here
Time to recovery of wall motion	Protocol-defined repeat imaging at fixed intervals, with a stated definition of normalisation	Not grounded here
In-hospital complications and mortality	Pre-specified complication set, admission to discharge, on a consecutive series	Not grounded here
Recurrence	Cumulative recurrence over a stated follow-up period, with defined ascertainment	Not grounded here
Depression and anxiety in this population	Case ascertainment by a stated diagnostic method, at a stated point after the acute event	**Qualitative statement of commonness only**

10.7 What the psychiatrist on the liaison call actually contributes

The consultation has a small number of jobs, and doing them well is the difference between a note that changes management and a note that restates the referral.

The first is to characterise the antecedent properly. A cardiology record says stress. A psychiatric record says what happened, at what hour, with whom, what it meant to this person, and how long the arousal lasted before the chest pain began. That timed narrative is the only evidence anyone will have about the trigger, it can be obtained on the day and rarely afterwards, and no other member of the team is trained to take it.

The second is to establish the exogenous and withdrawal load, using the differential above. This is a diagnostic contribution, because a withdrawal syndrome or a stimulant exposure changes what the cardiac team monitors and treats.

The third is to separate acute distress proportionate to a frightening event from a psychiatric disorder requiring treatment. The responses differ, and conflating them either medicalises normal fear or misses treatable illness.

The fourth is explanation, and it is underrated. The patient has been told they had something resembling a heart attack and also that their arteries are clean, and those messages conflict unless somebody supplies the mechanism reconciling them. A coherent, non-technical account of a surge, a stunned muscle and a recovering ventricle does real therapeutic work, replacing an incoherent threat with a comprehensible one.

The fifth is treatment and follow-up. **Where depression or anxiety is present, it is treatable in its own right, and a combined psycho-cardiologic rehabilitation approach is the model these patients might benefit from**. The obstacle is structural: a patient discharged with a normal angiogram is likely to leave with follow-up from neither service, and arranging it before discharge is the most useful line in the consultation note. Choice of agent in cardiac disease is a psychopharmacology matter and is handed over. Where consent arises, the operating principle is that capacity is decision-specific and time-specific, and the doctrine belongs to the mental health legislation notes.

RULE **Write the timed antecedent narrative on the day you are called.** Everything else in the consultation can be reconstructed later from records and from the patient. The hour-by-hour account of what happened before onset degrades within days, and it is the one datum that makes the case interpretable to anyone reading the notes afterwards.

IN INDIA

This build located no Indian series, registry entry, service description or cost datum for this syndrome, and none is supplied. What can be said structurally, without inventing a figure, is that recognition depends on early angiography, on repeat ventricular imaging to demonstrate recovery, and on a liaison service that receives the referral before discharge, and that the availability of each varies widely across Indian settings. A resident asked about the Indian picture is on firmer ground naming those dependencies and the absent data than quoting a prevalence figure with no Indian denominator.

GOOD TO REMEMBER

On the ward, hold these points. A clean angiogram excludes an obstruction and excludes nothing else. A psychological trigger never excludes a coronary occlusion. And the muscle is stunned rather than dead, which is why the story ends in recovery, and why the psychiatric work of characterising the antecedent, treating comorbid illness and arranging follow-up is done during an admission everyone else treats as resolved.

MNEMONIC

Surge, Switch, Stun, Settle

Surge is the catecholamine load from adrenal and neural sources. Switch is the apical receptor coupling change from stimulatory to inhibitory. Stun is living but non-contractile myocardium. Settle is the recovery that makes the diagnosis retrospective. The psychiatric work attaches at the start, on the antecedent, and at the end, on comorbidity and follow-up.

Beyond one territory
THE WALL-MOTION FEATURE THAT MAKES OCCLUSION IMPLAUSIBLE

Stunned
NOT INFARCTED, WHICH IS WHY IT RECOVERS

Common
DEPRESSION AND ANXIETY IN THIS POPULATION, STATED QUALITATIVELY

Timed narrative
THE DATUM ONLY THE PSYCHIATRIST CAN OBTAIN

A surge of catecholamines stuns living myocardium, the presentation copies an acute coronary syndrome exactly because the muscle really is injured, and the diagnosis is completed by recovery rather than by the angiogram. Take the timed antecedent on the day, look for exogenous and withdrawal catecholamine load, treat the comorbid depression and anxiety, and arrange the follow-up a normal angiogram will otherwise cancel.

Renal, transplant and perioperative psychiatry

11 · Pre-transplant psychiatric evaluation and post-transplant neuropsychiatry of immunosuppressants (steroids, tacrolimus, ciclosporin)

A transplant psychiatric opinion is read by a committee allocating a scarce organ, which makes it the one psychiatric assessment whose reasoning has to be legible to people outside

psychiatry. The topic splits cleanly and the sides are examined differently. Before transplantation, the question is what the psychosocial assessment must contain and how its conclusion should be written. Afterwards, it is which neuropsychiatric syndromes the immunosuppressants produce, and how to keep that possibility open while a recipient with several plausible causes of a new confusional state is worked up. The routing is strict, and one part of it is a gap rather than a destination. The general machinery by which a co-prescribed agent alters exposure, the inhibitor and inducer ladders and the cytochrome pathways, belongs to the psychopharmacology notes; that chapter names no immunosuppressant, so the interaction rankings and dose changes specific to these drugs are stated nowhere in this series and must not be reconstructed from memory. The confusional state itself belongs to the stroke, traumatic brain injury, delirium and focal brain syndromes notes. The steroid-induced psychiatric syndrome is routed on the plate below and taught in the endocrine and metabolic notes, while dialysis and the intensive-care sequelae around a transplant admission sit in the organ-interface routing section.

11.1 Where the opinion's form is owned, and what this section adds

Two limbs of this assessment are owned elsewhere and are named here rather than taught. How the opinion itself should be written, and the dual agency the interview runs under, belong to the Consultation-liaison, geriatric and transcultural psychiatry notes; so does the separate assessment of a living donor, where the question is whether the offer is voluntary and free of inducement, and where the psychiatrist works for the donor alone. The statutory framework governing donor authorisation belongs to the mental health and medical legislation notes. What this section adds is the limb those chapters do not carry: what the psychosocial assessment must contain, and how a capacity finding made during organ failure should be scoped, dated and re-asked.

11.2 The psychosocial assessment, and the limb this packet grounds

Guidance on evaluating candidates for kidney transplantation specifies the content of the psychosocial assessment, and the limb concerning psychiatry is stated explicitly. The elements of that assessment should include **a mental-status examination**, together with cognitive evaluation to ensure valid decision-making capacity and the ability to provide informed consent for transplantation.

Read that carefully, because the source is narrower than the way the topic is usually taught. It attaches the requirement to the psychosocial assessment as such, and attaches no timing qualifier of its own. This fragment therefore states it at that scope, and treats "pre-transplant" as a description of the setting in which the assessment is customarily performed rather than as part of what the guidance recommends.

The cognitive limb is the one candidates under-weight and the one most likely to change the answer. Organ failure impairs cognition by several routes at once, and the assessment often happens at the point of greatest impairment: uraemia, hepatic encephalopathy, hypoxia, poor cerebral perfusion and disrupted sleep. A patient examined in that state may appear unable to hold a decision they could hold comfortably once dialysis, oxygenation or ammonia handling

improved. So date the finding, name the reversible contributors, and set a point at which the question is asked again.

That licenses the one doctrinal point this section may state. Capacity is **decision-specific and time-specific**: a person may hold capacity for one decision and lack it for another on the same afternoon, and may regain it as their physiology improves. The statutory test, and the machinery that follows a finding of incapacity, are taught in the mental health legislation notes and named here rather than reproduced.

TRAP **Recording a capacity finding as a property of the patient rather than of a decision at a time.** In a population whose cognition tracks a treatable metabolic state hour by hour, writing "lacks capacity" as though it were a diagnosis removes a patient from consideration on a measurement taken at their worst, and the harm is invisible in the notes.

11.3 The rest of the interview, and where its evidence lives

Beyond the capacity limb, the assessment covers ground this packet does not ground with a source, so the honest move is to teach the structure of the interview rather than thresholds nobody here can cite. The domains a unit expects an opinion to address are the psychiatric history and current mental state; substance use and what any stated abstinence rests on; the adherence record; the caring arrangements after discharge; the patient's own account of the regimen and its permanence; and the financial arrangements that will keep immunosuppression running once the admission ends.

One interviewing discipline deserves stating outright, because it is where weak assessments go wrong. Ask about behaviour already on the record in preference to intention offered in the room. The record can be checked: dialysis attendance, missed appointments, prescription collection, what happened the last time a regimen became inconvenient. Why a stated intention is the weaker evidence here, given who the opinion is written for, is argued in the Consultation-liaison, geriatric and transcultural psychiatry notes. That is a point about where evidence lives, and this packet holds no figure licensing any claim that a given behaviour predicts a graft outcome.

Mental state and cognition	Record, not intention	Decision and time specific	Dated, then re-asked
WHAT THE ASSESSMENT MUST CONTAIN	WHERE ADHERENCE EVIDENCE LIVES	HOW CAPACITY IS SCOPED	WHAT A COGNITIVE FINDING IS

11.4 Tacrolimus, and the shape of calcineurin-inhibitor neurotoxicity

The calcineurin inhibitors are what make post-transplant care a psychiatric topic rather than a purely medical one. Their neurotoxicity is prominent enough in the labelling to belong on the differential of every neuropsychiatric presentation in a recipient. The tacrolimus prescribing information states the shape directly: tacrolimus may cause a spectrum of neurotoxicities, and the most severe of these include **posterior reversible encephalopathy syndrome (PRES), delirium, seizure and coma**.

Treat the word spectrum as the examinable part. A label describing a spectrum and then naming only its severe end is telling you a milder end exists while leaving it undefined. This packet holds

no row describing that milder end, so this fragment names the severe end the source names and goes no further.

What the statement licenses is a low threshold, and not a diagnosis. A recipient on tacrolimus who develops a seizure, a fluctuating confusional state or a depressed level of consciousness has a drug on the differential from the first minute, alongside infection, metabolic derangement, graft dysfunction and cerebrovascular events. The observation raising the drug's rank is temporal: whether the change followed an alteration in the regimen, a change of formulation, or the addition of another agent. Which agents do that, in which direction, and what dose response is then required are prescribing questions owned by the psychopharmacology notes.

Note the boundary inside that list: delirium appears as a manifestation of the drug's toxicity, which puts the drug into the delirium differential without putting delirium assessment into this section.

11.5 Ciclosporin, and a signal the label attaches to a combination

Ciclosporin is the other calcineurin inhibitor in routine use, and its labelling carries a neurotoxicity statement structured differently from the tacrolimus one, which is why it is worth learning separately. The statement is conditional on a co-prescription: convulsions have been reported in adult and paediatric patients receiving ciclosporin, particularly in combination with high-dose methylprednisolone.

Several features of that sentence repay attention.

- The signal is attached to a combination. A convulsion in the days after pulsed corticosteroid for suspected rejection is a combination event, and the history must establish what else was given and in what order, rather than settling on whichever agent the patient has taken longest.
- The labelling names adult and paediatric patients alike, so the statement spans the age range and a child or adolescent recipient attracts the same consideration.
- The wording reports occurrence rather than a rate. It licenses vigilance and a position on the differential; it licenses no statement about how often this happens.
- The named partner is a corticosteroid, so the same episode sits inside the corticosteroid material routed below. Naming the combination is this section's job; the steroid syndrome is not.

The plate below arranges the direct neuropsychiatric toxicity of the immunosuppressants by agent class rather than by syndrome, which is the arrangement answering the ward question: you know what the patient is taking and you want to know what it can do. The calcineurin inhibitors and the mTOR inhibitors are marked as interaction-sensitive, and that mark is a pointer rather than a teaching, since the interaction rankings and dose changes they oblige sit in the psychopharmacology notes. This packet holds no row describing an mTOR-inhibitor neuropsychiatric syndrome, so the prose asserts nothing about that class.

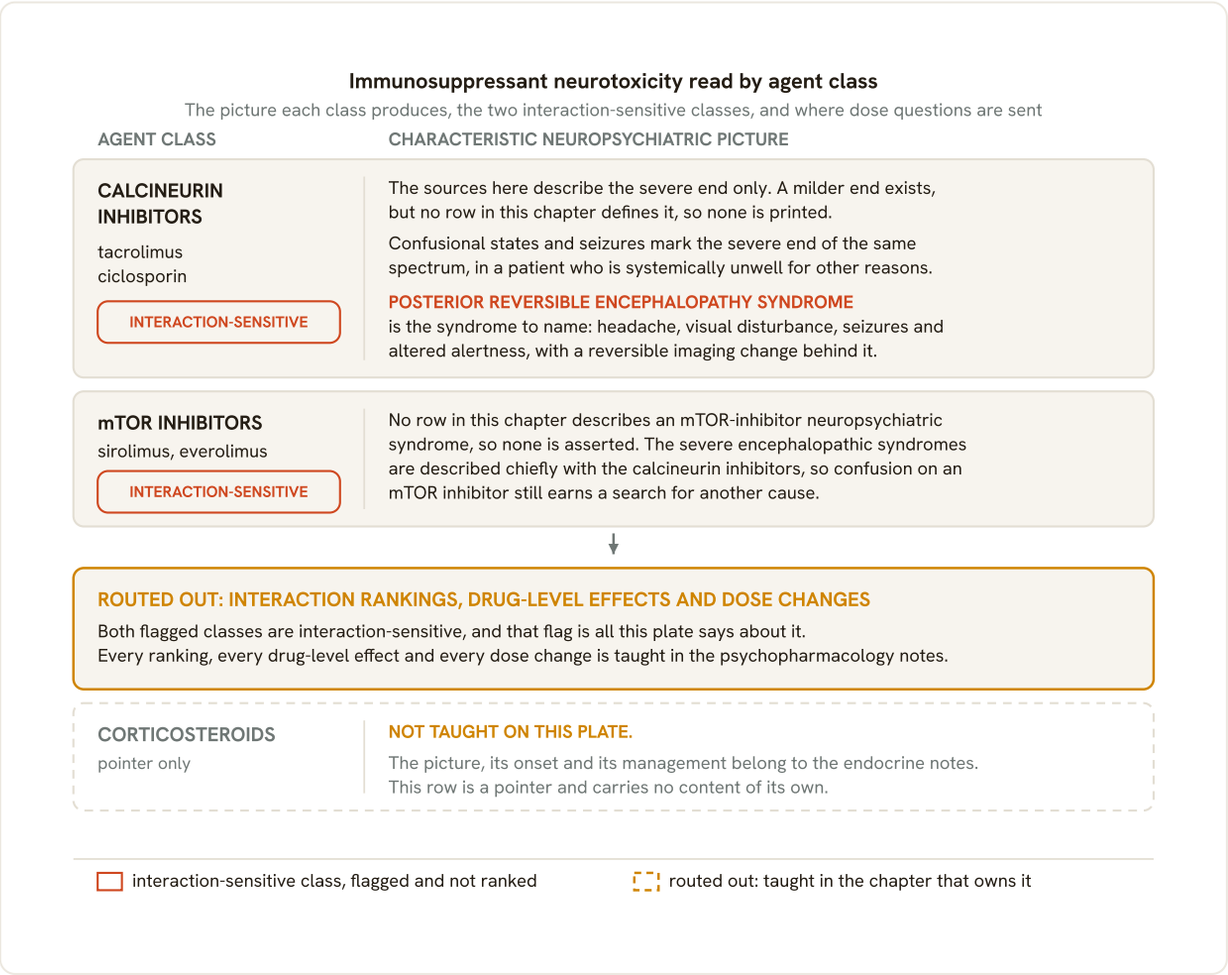


Fig 40.5 Direct neuropsychiatric toxicity of the immunosuppressants, read by agent class rather than by symptom. The calcineurin inhibitors carry the fullest picture: confusional states and seizures at the severe end, and posterior reversible encephalopathy syndrome as a named syndrome rather than a loose description of a confused transplant recipient. The milder end of that spectrum is left undefined here, deliberately, because no source in this chapter's evidence defines it, and naming a spectrum while supplying only one end of it from memory is how a plate acquires content its text does not hold. Nothing is asserted about an mTOR-inhibitor neuropsychiatric syndrome either, for the same reason; what the plate does say is that the severe encephalopathic syndromes are described chiefly with the calcineurin inhibitors, so confusion in a patient taking an mTOR inhibitor still earns a search for another cause before the drug is blamed. Three limbs are deliberately absent from the plate. Interaction rankings, drug-level effects and every dose change are taught in the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes; the corticosteroid picture is taught in the Endocrine and metabolic psychiatry notes; and the form of the pre-transplant opinion, together with the separate living-donor interview, belongs to the Consultation-liaison, geriatric and transcultural psychiatry notes. What that chapter records as a declared gap, the contents of the psychosocial assessment, is taught here. No Indian incidence figure is printed here, because the record was not located. From memory you should be able to name the class, give its severe picture, say the one syndrome out loud, and say which limbs this plate deliberately does not carry.

11.6 Corticosteroids, named here and taught elsewhere

Corticosteroids belong to the transplant regimen and therefore to this chapter, but the psychiatric syndrome they produce is owned elsewhere, so this section states the toxicity at label level and routes the clinical construct. Behavioural and mood changes sit among the common adverse reactions listed for corticosteroids, and that is the level at which this fragment may state them.

Steroid-induced psychiatric disturbance itself, meaning the mania, depression and psychosis that follow corticosteroid exposure and the management of each, belongs to the organ-interface

routing section of this chapter and to the endocrine and metabolic notes. An examiner asking about it in a recipient is asking a question whose answer is the same as in any other steroid-exposed patient: the transplant setting adds a co-prescribed calcineurin inhibitor and an immunosuppressed host, and adds no separate steroid syndrome.

11.7 The recipient with new neuropsychiatric signs, and the order of the questions

The characteristic difficulty after transplantation is that several sufficient causes are present at once, so the examinable skill is sequencing rather than selection. An immunosuppressed patient may be septic while barely febrile, metabolically deranged from graft dysfunction, hyponatraemic, hypoxic, and established on a calcineurin inhibitor and a corticosteroid at once, while carrying the psychological load of an operation they may have waited years for. An answer that names the drug and stops has selected the interesting cause rather than the first one.

The sequence that survives scrutiny puts reversible and dangerous causes first, keeps the immunosuppressant permanently on the list rather than at the end of it, and treats the psychological formulation as concurrent rather than as a diagnosis of exclusion. A recipient's distress about the donor, the graft, or a life reorganised around a regimen is real whether or not a metabolic cause is found, and deferring it until everything else is excluded means it is never addressed.

WHAT YOU OBSERVE	WHAT IT OBLIGES YOU TO ASK	WHERE IT IS TAUGHT
Fluctuating confusion or altered arousal	What is the medical cause, and what is the current regimen	Stroke, traumatic brain injury, delirium and focal brain syndromes notes
Seizure on a calcineurin inhibitor	What else was given, in what order, and when relative to the event	Label level here; prescribing in the psychopharmacology notes
Mania, depression or psychosis after corticosteroid exposure	Which agent, over what course, and what is the earlier psychiatric history	Organ-interface routing section; endocrine and metabolic notes
Doubt about a listing decision or an ongoing consent	Which decision, at what time, and what is reversible now	The requirement here; the statutory test in the legislation notes
Low mood on dialysis while awaiting a graft	What is the renal context, and what treatment is safe in it	Organ-interface routing section, which owns dialysis

Read the middle column rather than the first. Observation-to-diagnosis pairs fail on a stem where the observation has several sufficient causes; observation-to-question survives it, because the question stays the same whichever cause is operating.

11.8 The Indian setting, and what the assessment has to ask in it

Indian practice shapes this assessment more than any refinement of technique. Where the graft comes from a living relative, the psychiatrist acquires a second person to assess, with a different question and a different loyalty, inside a family in which pressure on a potential donor may be entirely unspoken. And because maintenance immunosuppression is a lifelong recurring cost, the sustainability of the regimen is a clinical variable rather than an administrative one.

IN INDIA

Ask the funding question in plain words, at the assessment and again before discharge: who will pay for immunosuppression once the admission is forgotten, and what happens if the recipient was the household's earner. Ask the travel question the same way, since a review schedule assuming a nearby clinic quietly excludes families who cannot make the journey. Where the donor is a relative, hold the separate donor interview the Consultation-liaison, geriatric and transcultural psychiatry notes describe, and record what was asked about voluntariness rather than only the conclusion. This packet holds no Indian figure for transplant volumes, living-donor proportion, drug cost, or psychiatric morbidity in Indian cohorts, so this box states the questions and states no magnitude.

GOOD TO REMEMBER

A short set of sentences carries most of the marks. The psychosocial assessment should include a mental-status examination and cognitive evaluation of decision-making capacity and the ability to provide informed consent for transplantation. Tacrolimus produces a spectrum of neurotoxicity whose severe end includes PRES, delirium, seizure and coma. Convulsions have been reported with ciclosporin, particularly in combination with high-dose methylprednisolone, and behavioural and mood changes are listed among the common corticosteroid adverse reactions.

PITFALL

Attributing a post-transplant seizure to ciclosporin alone. The labelled signal is conditional on a combination with high-dose methylprednisolone, so an answer naming the calcineurin inhibitor and omitting the corticosteroid has reported part of the statement and has stopped taking the history where it was about to become informative. The reverse error costs the same marks.

The pre-transplant side this section owns is the content of the psychosocial assessment, a mental-status examination and a dated cognitive evaluation of capacity, with the opinion's own form owned elsewhere; the post-transplant side is a permanently open differential in which tacrolimus carries a spectrum reaching PRES, delirium, seizure and coma, ciclosporin carries a convulsion signal tied to high-dose methylprednisolone, and corticosteroids are named here and taught elsewhere.

12 · Perioperative psychiatric management, capacity for surgery and postoperative cognitive dysfunction

Surgery reliably produces a new psychiatric presentation in patients who had none the week before. Perioperative psychiatry is best learned as a sequence, because the questions arrive in a fixed order and each has a different owner. Before the operation the question is optimisation and measurement. At consent it is capacity. Afterwards it is which cognitive construct is being named. This section installs that sequence and makes the hand-offs a candidate is expected to make rather than answer: delirium assessment and management, taught with the stroke, brain injury and focal syndrome material, and capacity doctrine with its statutory test, taught with the mental health legislation material.

12.1 The perioperative sequence, and who owns each question

A psychiatrist meets a surgical patient in a small number of recognisable ways, and naming them is the fastest route to a structured answer. The first is the planned preoperative review of a patient with established psychiatric illness, asking whether the illness or its treatment changes the operation or the recovery. The second is the urgent review at consent, where the team is unsure whether the patient can make the decision in front of them. The third, by far the commonest, is the postoperative call about a patient who was lucid on admission and is now confused or agitated. The fourth, most often missed, is the late review months afterwards, where the complaint is not confusion but persisting difficulty with memory or work.

Those shapes belong to different literatures and different chapters, and an examiner wants them separated before the answer starts. The characteristic failure is a single undifferentiated account of postoperative confusion, which loses marks in both directions: it imports delirium teaching into a question about cognitive trajectory, and trajectory language into a question about a disturbed patient on the ward tonight. So fix ownership first. The acute confusional presentation is the delirium chapter's, entire; consent and its statutory machinery are the legislation chapter's; agent, dose and organ-adjusted prescribing are the psychopharmacology chapter's. The plate below traces the pathway from preoperative review, through the operation, to the postoperative neurocognitive outcome, marking each exit where it occurs.

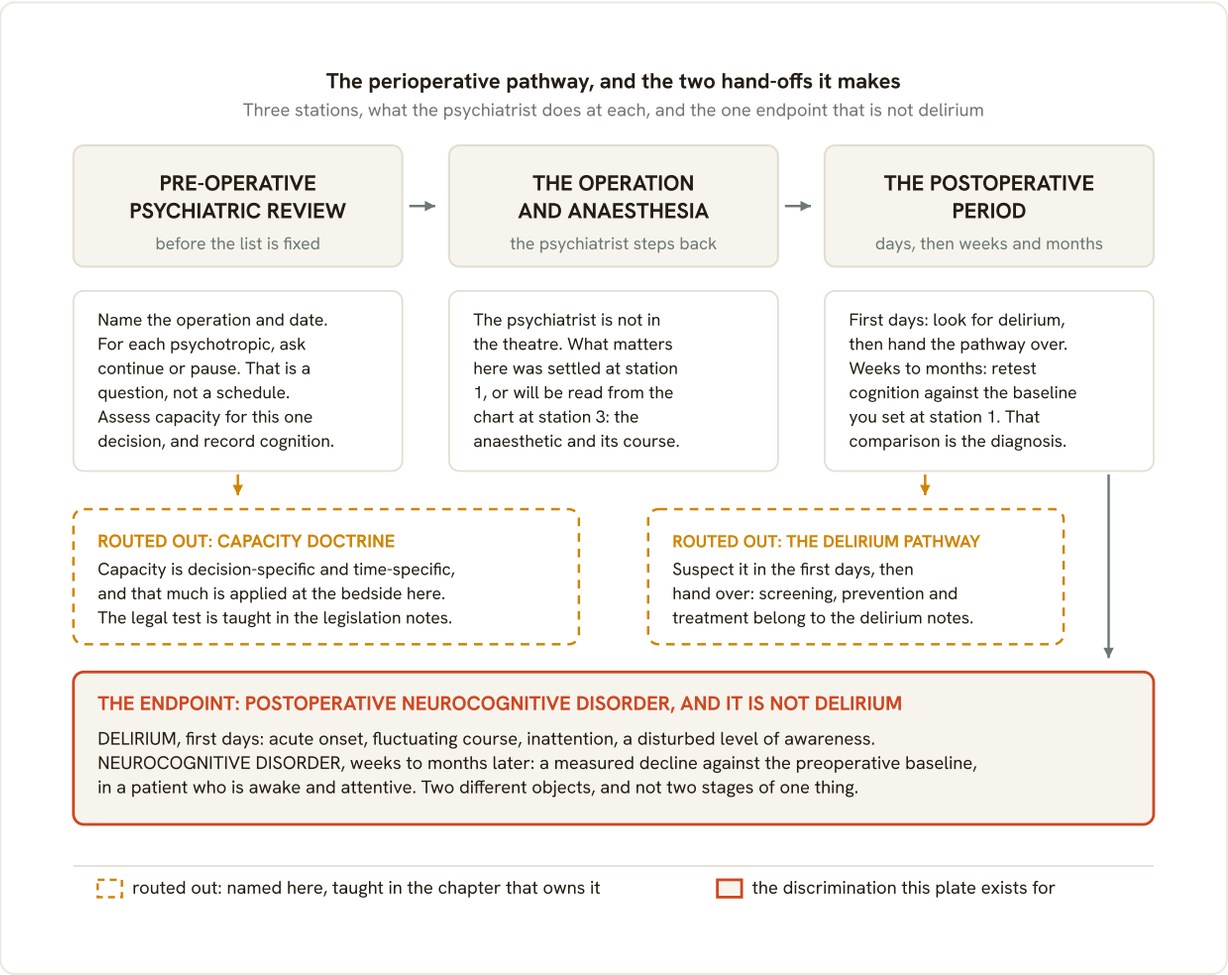


Fig 40.6 The perioperative pathway drawn left to right, with the psychiatrist's task at each station and the two hand-offs marked as branches off the rail. Station 1 is where the work is actually done: medication decisions are framed as questions rather than as a schedule, capacity is assessed for this one decision at this one time, and the cognitive baseline is recorded, because without that baseline the endpoint at station 3 cannot be measured at all. The capacity branch routes to the Mental health legislation and forensic psychiatry notes, which hold the legal test, and consent to the surgical decision is set out in the Forensic ethics, consent and legal procedure notes; the delirium branch routes to the Stroke, TBI, delirium and focal brain syndromes notes, which hold screening, prevention and treatment. The endpoint is postoperative neurocognitive disorder, and a nomenclature working group has recommended perioperative neurocognitive disorders as the overarching term for cognitive impairment identified in the preoperative or postoperative period. From memory you should be able to state what the psychiatrist does at each station, name both hand-offs, and separate delirium from postoperative neurocognitive disorder on onset, attention and course.

12.2 The name of the thing, and why the name was changed

The older vocabulary here was built by researchers rather than clinicians, and it showed. Terms for cognitive change after anaesthesia and surgery accumulated separately, each defined by whichever measurement a study happened to use, so a single patient could be described in incompatible ways depending on which literature the describer had read.

A nomenclature consensus working group addressed this by proposing a single umbrella. **The working group recommends that perioperative neurocognitive disorders be used as an overarching term for cognitive impairment identified in the preoperative or postoperative period.** Read that formulation slowly, because it does more than rename an old idea.

Notice what kind of statement it is: a recommendation about terminology, issued by a working group, and it should be presented as such. An answer promoting it into a diagnostic category with

criteria and thresholds has claimed more than the recommendation says. Notice next where the umbrella reaches. It covers impairment identified in the preoperative period as well as the postoperative period, which is the counter-intuitive limb and the examinable one. That decision converts preoperative cognitive assessment from a chore attached to the construct into part of the construct itself.

RULE **The umbrella reaches backwards, not only forwards.** The overarching term covers cognitive impairment identified before the operation as well as after it. A candidate who describes it as a class of postoperative complication has described half of it, and has lost the reason preoperative assessment belongs in the same discussion.

12.3 The preoperative psychiatric review, and what it is actually for

A preoperative review earns its place by settling a short list of questions where the surgical and anaesthetic teams will find the answers. The first is continuity of psychotropic treatment across the fasting window and any period in which the patient cannot swallow. The review does not choose the agent, which belongs with every dose and interaction ranking to the psychopharmacology material, but it must state, for each agent, whether it continues, is withheld or is given by another route, and which team owns that decision. A plan recording only that the patient takes long-term medication has settled nothing.

The second is withdrawal. An elective admission removes access to alcohol, to nicotine and to anything else the patient used at home, at a predictable hour. A withdrawal syndrome declaring itself on a surgical ward is read, by people looking for a surgical explanation, as postoperative confusion. Establishing what the patient used before admission is preoperative work, and recording it where the postoperative team will see it is the same job.

The third is the illness features that change postoperative risk without changing the operation: patients who communicate pain poorly, who do not seek help when something is wrong, who become agitated in unfamiliar surroundings, or whose treatment history makes an interruption expensive. Each is a practical problem invisible to a clerking that asks about organs. Pre-transplant psychiatric evaluation shares this shape with a different content and a different threshold, and is taught by the transplant section of this chapter.

12.4 The preoperative cognitive baseline, and why it is not optional

The most useful contribution a preoperative review makes to the postoperative period is a measurement taken before the operation. **Older patients should have preoperative cognitive assessment using established screening or diagnostic tools.** The recommendation names a class of instrument rather than a particular one, and this section follows it exactly. Naming a specific instrument here would need a source that attaches a property to it, and none is held, so supplying one from recollection is the error this chapter exists to prevent.

The reasoning is worth stating, because it is what makes the recommendation defensible rather than merely official. Postoperative cognitive change is a change, and a change is a relation between two observations; one postoperative observation cannot establish it. Without a

preoperative measure the team compares the patient against an impression of the patient, formed by relatives under stress and reported after the event.

There is a quieter second yield. Impairment found at the preoperative assessment is informative before any operation happens: it marks a patient whose recovery needs more support, it bears on how the consent conversation is conducted, and it frequently identifies impairment nobody had recognised.

TRAP **Reporting a low postoperative score as a postoperative deficit.** Candidates do this because a score looks like evidence and because the postoperative score is the one they have. Without a preoperative measure it establishes the patient's current level and nothing about change, and a patient with unrecognised pre-existing impairment generates the same result. The examinable point is not the score, it is the comparison the score was supposed to enter.

12.5 Capacity for surgery, and the part of it that belongs here

Almost all of capacity doctrine sits elsewhere. The statutory test and its limbs, the presumption the law starts from, the safeguards and the procedure that follows a finding of incapacity are taught with the mental health legislation material, and none of it is restated here or cited by provision. One property does belong in this section, because the perioperative period is where it does most of its work. **Capacity is described as decision-specific and time-specific.**

Decision-specific means the question is never whether the patient has capacity as a standing attribute, but whether this patient can make this decision. A patient may be unable to manage a complex financial arrangement and entirely able to decide about an operation, and the reverse holds equally. An opinion recorded for one purpose therefore does not answer a different question asked later.

Time-specific means the answer attaches to a moment rather than to a person, and the perioperative period supplies an unusual density of moments in which it can move: pain, disturbed sleep, premedication and the hours immediately after anaesthesia all sit inside the window in which consent is sought. An opinion formed on the morning ward round is not transferable to a consent conversation held that evening. What the psychiatrist supplies is a description of the patient's functional abilities for the specified decision at the specified time, in the terms the legislation sets out; the decision itself belongs to the treating team.

PITFALL

Recording a global verdict. A note stating that the patient lacks capacity, with no decision named and no time attached, is unusable by the team receiving it and indefensible if it is later examined. It also travels: written once, it is quoted for months, across decisions it was never about. Name the decision, name the time, confine the opinion to both.

12.6 Postoperative neurocognitive disorder distinguished from delirium

The two constructs are distinguished by what kind of thing each one is, not by comparing symptom lists. Delirium names an acute disturbance of attention and awareness with a characteristically fluctuating course. It is a condition in its own right, with its own screening

instruments, precipitant list and management, all taught with the stroke, traumatic brain injury and focal brain syndrome material and none of it repeated here. Postoperative neurocognitive disorder names something structurally different: a change in cognitive performance relative to a preoperative baseline, established by comparison rather than by observation at the bedside.

That difference in kind produces every other difference an examiner will want. One is recognised by looking at the patient now; the other cannot be, because a single observation carries no baseline. One fluctuates, which is why repeated observation detects it; the other is a persisting shift, which comparison rather than repetition finds. One declares itself on the ward; the other declares itself after discharge, and reaches a clinician only if somebody thinks to ask.

THE QUESTION	DELIRIUM	POSTOPERATIVE NEUROCOGNITIVE DISORDER
What the term names	An acute disturbance of attention and awareness	A change in cognitive performance against a preoperative baseline
How it is established	Observation of the patient's current state	Comparison of a later measure against an earlier one
Characteristic course	Fluctuating, often worse at night	Persisting, without the same fluctuation
Where it declares itself	On the ward, to the clinical team	After discharge, to the patient and family
Who teaches assessment and management	The stroke, brain injury and focal syndrome material	This section for the construct; psychopharmacology for any agent

RULE **Delirium is observed; a perioperative neurocognitive disorder is computed.** One is a state visible in the room, and the whole of its assessment lives in the chapter that owns it. The other is a difference between two measurements, which is why a service that never takes the first measurement can never make the diagnosis, however carefully it examines the patient afterwards.

The recommendation governing the acute limb is stated here only as a boundary. **Postoperative delirium screening should use a validated delirium score.** Which score, how it is administered and what follows a positive result are the delirium chapter's, and a candidate who reproduces an instrument here has answered another chapter's question while believing they were thorough. The perioperative point is narrower: screening is systematic and instrumented rather than impressionistic, which is a statement about how a surgical service is organised.

The constructs are not alternatives to choose between, and the same patient may pass through both, in sequence, after one operation. Whether an episode of postoperative delirium predicts the later cognitive trajectory is a question no source held for this section answers, and it is recorded below as a gap rather than taught.

MNEMONIC

Measure before, ask at, name after

Measure cognition before the operation, because a change needs an earlier observation. Ask the capacity question at the decision and at the moment. Name the construct afterwards, because delirium and a perioperative neurocognitive disorder are different kinds of thing with different owners.

12.7 Answering the postoperative call without importing the wrong chapter

The postoperative call is where most candidates meet this material, and the discipline that rescues an answer is to establish the shape of the problem before naming it. Establish whether a preoperative cognitive measure exists, because that determines what can be concluded at all. Establish the timeline, separating a disturbance that arrived over hours from a difficulty the family has watched for weeks. Establish whether the disturbance is of arousal and attention now, or of performance over time. Establish what else on the chart could produce it, since a surgical patient carries pain, immobility, infection risk and an interrupted medication list at once.

A postoperative note worth reading later states each of the following.

- Whether a preoperative cognitive assessment exists, where it is recorded, and what it showed.
- The timeline of the change, expressed against the operation rather than against the referral.
- Which construct is being described, and which is being excluded.
- The medication position: what was withheld, what was resumed, what is outstanding.
- Who reviews the patient after discharge, since the trajectory question cannot be answered on the ward.

Routing follows the same discipline. Psychological sequelae of a critical-illness stay, and delirium at the end of life, are taught by the organ-interface section of this chapter. Every question about agent, dose and organ failure belongs to the psychopharmacology material. Naming the destination answers that part of a question completely.

Before it, and after it	Measured, not recalled	This decision, this moment	Observed, or computed
WHAT THE UMBRELLA TERM COVERS	WHAT A LATER SCORE IS COMPARED WITH	HOW THE CAPACITY OPINION IS SCOPED	WHICH CONSTRUCT IS BEING NAMED

Every tile scopes something rather than reporting it: the reach of the term, the comparison a later score must enter, the boundaries of a capacity opinion, and which of the two constructs is in play.

12.8 The Indian perioperative reality, and what it changes

The perioperative recommendations held here come from a European perioperative safety guide, and the nomenclature recommendation from an international consensus working group. No Indian perioperative cognitive-assessment standard, and no Indian figure for how often preoperative cognitive assessment is performed, is held by any source behind this section. That absence is stated rather than filled, and it should be stated in an answer too.

What transfers is clear enough to reason about. The preoperative cognitive baseline is among the cheapest recommendations in perioperative medicine: a brief structured assessment, a person to administer it, and a place in the record the postoperative team will find. The barrier in most Indian surgical services is not cost but ownership, because preoperative clerking belongs to the surgical and anaesthetic teams, and the cognitive item is the first dropped when the list is long. Language and literacy add a second consideration: an instrument administered in a language the patient does not use, or whose items assume schooling the patient did not have, produces a low score that is a finding about the instrument rather than the patient.

IN INDIA

The practical move for an Indian unit is to attach the cognitive item to the anaesthetic preoperative assessment, where older surgical patients are already seen and documented, rather than to build a separate psychiatric pathway that depends on a referral being made. The psychiatric contribution is then to specify which instrument the unit uses, in which languages, administered by whom and recorded where, and to accept that a brief assessment applied consistently beats an ideal one done occasionally. No Indian coverage or outcome figure for such a pathway is held here.

GOOD TO REMEMBER

When the surgical team telephones about a confused postoperative patient, the first question is not what is wrong but whether anybody measured them before. If yes, a genuine comparison is available and a trajectory question can be answered. If no, say so plainly in the note, manage the acute presentation through the chapter that owns it, and arrange the review after discharge.

Delirium is seen at the bedside and belongs to another chapter; a perioperative neurocognitive disorder is a comparison against a preoperative baseline, which is why the measurement taken before the operation is this section's whole practical yield.

Oncology, palliative and paraneoplastic psychiatry

13 · Psychiatric effects of corticosteroids, chemotherapy and interferon (steroid psychosis, chemobrain)

The examinable content in cancer-drug neuropsychiatry is not a list of side-effects but a mapping: which agent, which syndrome, and over what interval that syndrome appears and is looked for. A candidate who has memorised that cytotoxic chemotherapy is associated with cognitive complaints, and that interferon-alpha is associated with depression, has learned the two least discriminating facts available. Marks sit on the rest: where in the course a syndrome belongs, what else produces the same picture at that point, whether the complaint has an objective correlate, and what the evidence licenses in advance rather than in response. This section installs that mapping for two agent families and deliberately routes a third away.

13.1 The boundaries this section keeps

Three boundaries govern what follows. Each exists because the material sits in full elsewhere in these notes, and a second version written here would drift from the first without anyone noticing.

The corticosteroid limb of this topic is not taught here. Systemic corticosteroids and steroid-induced psychiatric disturbance are routed on the chapter's opening plate to the endocrine and metabolic psychiatry material, and the merged organ-interface routing section of these notes carries the clinical construct. This section names steroids and goes no further with them.

Organ-adjusted prescribing is elsewhere too. Doses, dose adjustments and interaction rankings between psychotropics and cytotoxic agents belong to the psychopharmacology notes, and no figure of that kind appears below. The commonest way a candidate loses control of this topic is to answer a question about syndrome recognition with a half-remembered interaction table.

The third boundary is interferon. **Interferon-alpha** is taught here. **Interferon-beta** is taught in the Neuropsychiatry of systemic and neurological disease notes. That is a boundary rather than an accident of filing, and a later subtopic gives the reason.

RULE **Name the agent, name the syndrome, name the interval.** An answer that pairs an agent with a syndrome and stops has given half of what a ward needs. The interval tells the clinician when to look and which explanations compete at that moment.

13.2 Cancer-related cognitive impairment, and why the patient's word is the wrong construct

Patients say **chemobrain**. The construct a clinician should hold is **cancer-related cognitive impairment**, and the change of name is not fastidiousness. A term that asserts its own cause stops the clinician looking, and the value of this topic lies in what is still worth looking for after the complaint has been taken seriously.

The presenting complaint is characteristically about effortful cognition rather than frank forgetting: holding several things in mind, sustaining attention through a task, retrieving an available word, working at the old speed. Patients describe the change against their own baseline rather than any norm, which is both why the complaint is genuine and why it is hard to demonstrate.

Note what the study design in the next subtopic concedes before a single result is read. Its comparison groups are patients with the same cancers, not healthy volunteers. A design that needs disease controls is admitting that the disease, its investigations, its surgery, its other treatments and the fatigue around all of them are candidate causes of the same complaint. The cytotoxic agent is one contributor inside a bundle, and the patient's word names the one easiest to blame.

That study measured cognition with self-reported and performance-based cognitive measures, and it is worth registering that it used both. A complaint and a test result are different observations: the first cannot tell you whether performance has changed, and the second cannot tell you whether the person's life has changed.

13.3 The longitudinal evidence, read for exactly what it licenses

The anchor for the time course here is a longitudinal study of patients receiving chemotherapy for two tumour types, each type carrying its own control group. **225 participants** were assessed. The breast-cancer stratum held **41 patients and 36 controls**; the lymphoma stratum held **73 patients and 75 controls**. The reported conclusion is that, compared with controls, patients with breast cancer and lymphoma showed persistent cognitive worsening still present at **1 year** after treatment.

Read the boundary of that sentence carefully, because it is where the marks are. It is a group-level comparison: it describes two treated groups against two disease-matched groups and predicts nothing about the individual in front of you. It covers two tumour types, so it says nothing about regimens used for other cancers. The strata are modest, and the lymphoma stratum is the larger, so generalise more slowly from the smaller.

TRAP **Converting an observation window into a prognosis.** A finding reported up to a stated point after treatment describes how far the looking reached, not where the problem ended. Candidates translate it into reassurance, that it settles within the year, or into catastrophe, that it is permanent. The evidence supports neither, only that worsening was still detectable at the end of the window examined.

The consequence of persistence is a scheduling consequence. If the complaint outlasts the treatment, the point at which oncology hands the patient back is the point at which nobody is reviewing cognition, and it is exactly then that the patient is trying to return to work. A service that reviews cognition only while chemotherapy runs has arranged to stop looking just before the problem becomes disabling.

13.4 Assessing the complaint without reaching for an instrument

The task is not to confirm the patient's attribution. It is to characterise the complaint, subtract everything treatable that produces the same picture, and decide whether anything acute or evolving sits underneath it. The order matters: the treatable contributors are common and the acute ones are dangerous, and both are missed by a clinician who has already accepted the diagnosis the patient offered.

THE QUESTION TO ASK	WHY IT DISCRIMINATES	WHERE THE ANSWER ROUTES
Did the difficulty predate the first cycle?	Separates a complaint attributed to treatment from one already present at diagnosis	Keeps the differential open instead of closing it on the drug
Is the picture stable across the day, or fluctuating with clouded attention?	Fluctuating inattention is an acute-state question, not a chronic-impairment one	The stroke, traumatic brain injury, delirium and focal brain syndromes notes
Is mood, sleep, pain or fatigue carrying the complaint?	Effortful cognition fails when drive and sleep fail, and those are addressable now	The depression, demoralisation and adjustment material in the merged organ-interface routing section
Is the picture static, or evolving between reviews?	An evolving encephalopathy is a different class of problem from a stable deficit	The paraneoplastic neuropsychiatric syndromes section; the workup sits in the neurology-psychiatry interface notes
Does the complaint have an objective correlate on testing?	Self-report and performance are separate observations answering different questions	Decides whether the conversation is about workplace accommodation or further assessment

Two further routings belong in the same breath. Capacity, where a treatment decision turns on it, is decision-specific and time-specific, and the statutory test is taught in the mental health legislation notes. Where the complaint sits alongside a request to stop treatment, the palliative and existential material in the merged organ-interface routing section owns that conversation.

GOOD TO REMEMBER

The most useful thing a liaison psychiatrist adds here is a timeline. Ask when the difficulty began relative to diagnosis, to surgery, to the first cycle and to the end of treatment, and ask which task first went wrong. A complaint that predates any cytotoxic agent, one that fluctuates within a day, and one still worsening after treatment finished are three different problems wearing the same words.

13.5 Interferon-alpha, and the one thing this section can state as evidence

A spelling point first, because it costs marks in reading rather than in writing. The oncological preparation is written **interferon alfa** in drug nomenclature and in much of the trial literature, and interferon-alpha in general prose. They are the same agent.

The evidence this section can state concerns prevention rather than treatment. In patients with malignant melanoma who were eligible for high-dose interferon alfa therapy, pretreatment with paroxetine appears to be an effective strategy for minimising depression induced by interferon alfa. That was tested in a randomised comparison in which **40 patients were randomised, 20 to paroxetine and 20 to placebo**, with symptoms consistent with major depression during interferon alfa therapy as the outcome of interest.

Retain the qualification exactly as the source states it. Pretreatment appears to be an effective strategy, which is weaker than a claim that it prevents the syndrome, and an answer that upgrades

it has misreported the source. The population is equally load-bearing: melanoma, high-dose interferon alfa, one antidepressant, and pretreatment given before the interferon rather than treatment started once mood has changed. None of those boundaries transfers on its own to another cancer, another interferon, another antidepressant or a reactive prescription.

One thing can be inferred honestly from the design, and it carries the time course. A placebo-controlled prevention trial is only designable if the syndrome it prevents is expected often enough, and soon enough after therapy starts, to be counted inside the treatment course. The trial's shape therefore tells you that interferon-alfa depression is foreseeable and time-linked rather than idiosyncratic. That is an inference about the design, and this section publishes no rate.

RULE **Prophylaxis is a different claim from treatment, and a population is part of a claim.** The bound evidence supports pre-emptive antidepressant cover in a specified oncology population starting a specified interferon regimen. It says nothing about what to do once the mood syndrome has declared itself, and it does not extend to everyone about to receive an interferon. Quoting a prevention result as a treatment result is an invisible error unless you state the population every time you state the finding.

13.6 Interferon-alpha is not interferon-beta, and the separation is clinical

These are distinct preparations, given to distinct populations for distinct reasons, and their psychiatric literatures are separate. Interferon-alpha reaches the psychiatrist through oncology and through infectious disease, in courses that are intensive and time-limited. Interferon-beta reaches the psychiatrist through a chronic neurological illness carrying its own mood burden independently of any drug, which is why the Neuropsychiatry of systemic and neurological disease notes hold it and this section does not.

The separation matters for three reasons an examiner can ask about directly. First, the background rate of psychiatric disorder differs between the populations, so a figure pooled across both would describe no real patient and mislead in both directions at once. Second, the confounders differ: in the alpha setting the competing explanation is the cancer or the infection and its treatment, in the beta setting the neurological illness itself. Third, the clinical decision differs. In the alpha setting the live question is whether to cover a foreseeable, time-linked mood syndrome before a defined course begins; in the beta setting it is embedded in long-term management of a chronic disease, and taken up where that disease is taught.

TRAP **Answering "the neuropsychiatric effects of interferon" as a single topic.** A stem that says interferon without saying which preparation is testing whether the candidate knows the distinction exists. The answer that scores separates alpha from beta, states which one it will discuss, and says why pooling them would be uninformative. The answer that fails treats the word as one drug and produces a paragraph true of nobody.

13.7 Agent to syndrome to time course, which is what the plate carries

The plate below is the compressed form of everything above, built on three columns rather than two. The first names the agent family. The second names the syndrome, in the vocabulary a

clinician would use in a note. The third, which does the work, places that syndrome on the treatment timeline: before treatment, during the course, at handover, and after treatment has finished. Read the third column first when revising, because it converts a list of associations into something usable at the bedside.

Two features of the plate are deliberate rather than accidental. The corticosteroid limb is absent, because steroids are routed on the chapter's opening plate to the material that owns them. And interferon-alpha and interferon-beta are drawn apart rather than sharing a row, for the reasons given above.

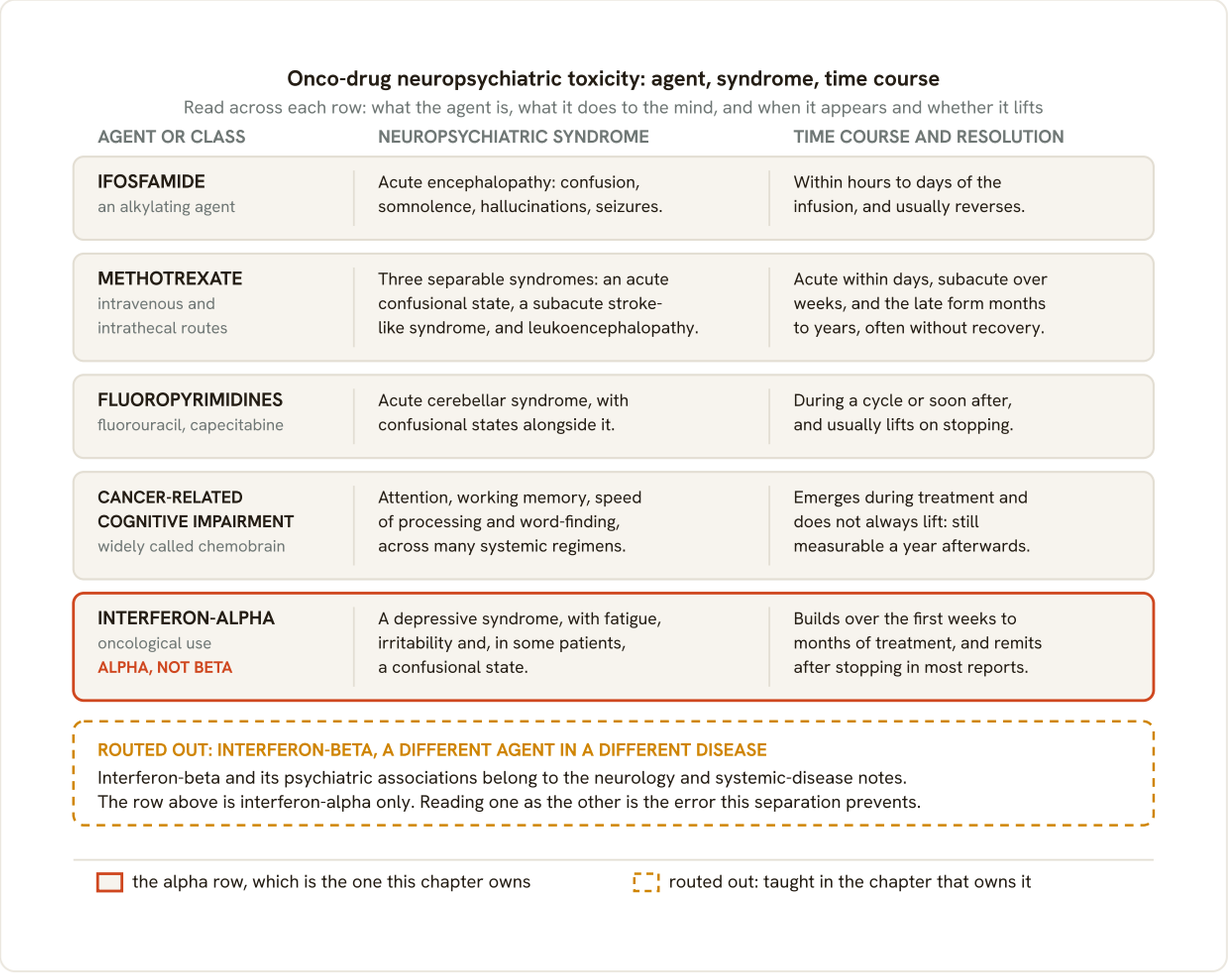


Fig 40.7 Chemotherapy and interferon neuropsychiatric toxicity read as agent, then syndrome, then time course, because at the bedside it is the time course that usually settles the differential. The acute encephalopathies declare themselves within hours to days and usually reverse on withdrawal, which is why ifosfamide, the fluoropyrimidines and the acute methotrexate syndrome sit together; methotrexate is the agent that also carries a subacute stroke-like syndrome and a late leukoencephalopathy that may never recover. Cancer-related cognitive impairment, widely called chemobrain, is the slow limb: it emerges during treatment, spares alertness, and has been measured as still present a year after treatment ends. Interferon-alpha produces a depressive syndrome that builds over the first weeks to months of treatment. The plate separates interferon-alpha from interferon-beta deliberately, because they are different agents used in different diseases and beta belongs to the Neuropsychiatry of systemic and neurological disease notes; the third drug group named in this topic is carried by the routing plate earlier in this chapter and by the Endocrine and metabolic psychiatry notes, so it is not drawn here at all. From memory you should be able to place any of these agents on the three columns and say whether the syndrome is expected to lift.

IN INDIA

No Indian figure for the frequency of cancer-related cognitive impairment, and none for interferon-alfa depression in Indian oncology practice, was located for this section, and none is asserted. What transfers is the structure. Where oncology and psychiatry sit in different institutions and follow-up is episodic, the practical question is not which cognitive instrument to buy but who reviews the patient once the oncology course closes, and when. Since the bound evidence describes a complaint still present after treatment ended, a service with no review point after discharge from oncology has designed itself to miss the window that matters most.

13.8 The shape of an answer that scores

A written answer here is graded on discrimination rather than length, and the discriminating content is ordered.

- Separate the agent families at the outset. Name the corticosteroid limb only to route it, and separate interferon-alpha from interferon-beta before saying anything about either.
- Use the construct rather than the patient's word, and say why the patient's word closes the differential prematurely.
- State every finding with its population attached. A finding without its population is not a shorter answer, it is a weaker and different one.
- Place each syndrome on the treatment timeline, and keep the observation window distinct from the prognosis.
- Route the neighbouring problems by name: delirium, paraneoplastic syndromes, depression and adjustment in cancer, capacity, and organ-adjusted prescribing.
- Say what is not known. The honest gaps here are large, and a candidate who names them is more convincing than one who fills them.

PITFALL

Treating the cognitive complaint as the end of the assessment because the patient has a plausible reason for it. A patient receiving chemotherapy has an explanation available for almost any neuropsychiatric presentation, and that availability is what makes delirium, an evolving encephalopathy and an untreated depression easy to miss. The plausible cause is the one to test last, not the one to accept first.

Agent, syndrome, interval
THE THREE THINGS AN ANSWER MUST NAME

Cancer-related cognitive impairment
THE CONSTRUCT THAT REPLACES THE PATIENT'S WORD

Prevention, not treatment
WHAT THE PRETREATMENT FINDING LICENSES

Alpha is not beta
THE SEPARATION A ONE-WORD STEM IS TESTING

Interferon-alpha is separated from interferon-beta and corticosteroids are routed away; the cognitive syndrome is cancer-related, not chemotherapy-proven; and every finding is stated with its population and its observation window attached.

14 · Paraneoplastic neuropsychiatric syndromes and onconeural antibodies: when to suspect occult malignancy

A paraneoplastic neuropsychiatric syndrome is a psychiatric presentation driven by the immune response to a tumour that has not yet been found. The final clause carries the whole examinable weight. This is not a complication of a cancer already sitting in the case notes; it is the event that makes someone go looking for the cancer in the first place. This section installs the reasoning that converts an ordinary-looking psychiatric presentation into a directed search for occult malignancy: what the term **paraneoplastic** actually claims, why the psychiatric clinic meets these patients ahead of the oncology clinic, what raises suspicion, what a psychiatrist owns and what is handed on, and which antibody has been reported to point where.

14.1 What the term claims, and the sequence that defines it

The prefix does the work. **Para** plus **neoplastic** names a disorder that accompanies a neoplasm while being made of something else. The damaged tissue is nervous tissue, and what damages it is an immune response the tumour provoked. The word **onconeural** names why such a response reaches the brain at all: the antigen belongs to the tumour and to neural tissue alike, so an attack aimed at one compartment lands in the other. Everything examinable here follows from that idea, including why the syndrome can be severe while the tumour stays small enough to have escaped detection.

The defining feature, for examination purposes, is a sequence rather than a severity. **The neurological or psychiatric syndrome comes first, and the malignancy is detected afterwards.** The task-force report that grounds this section states it plainly: **paraneoplastic neurological syndromes almost invariably predate detection of the malignancy**. Hold that sentence in mind through everything that follows, because it is what separates this topic from every other cancer topic on the day, and it is the reason the topic belongs to a psychiatry paper at all.

It is worth being precise about what the label has to displace before it is earned, because a neuropsychiatric picture in someone who turns out to have cancer has several ordinary explanations. It may be a direct effect of tumour mass or of secondary deposits. It may be metabolic or endocrine derangement produced by the tumour. It may be treatment: the cytotoxic agents, the interferon and the corticosteroids, each taught in its own place, in the sibling section on chemotherapy and interferon effects and in the routing section that owns steroid-induced disturbance. It may be adjustment, demoralisation or depression in someone absorbing a serious illness, which the routing section on cancer owns. Cross-refer each by name; this section restates none of their content.

Notice the asymmetry the sequence creates. In the ordinary cancer psychiatry of a known malignancy, those rival explanations are all available on the ward round, because the tumour, the metabolic panel and the drug chart are already in front of you. In the paraneoplastic case they are unavailable at presentation, because nobody has yet established that there is a tumour. The

clinician holds a syndrome and a history, and the question of malignancy has to be raised rather than handed over.

RULE **The defining feature is the sequence, not the severity.** **A paraneoplastic syndrome characteristically declares itself before the malignancy is detected**, which is why the phrase in the bank item is "when to suspect occult malignancy" and why an answer that describes a florid syndrome in a known cancer patient has answered a different question.

14.2 Why the psychiatric clinic meets these patients first

There are honest structural reasons why this presentation lands with a psychiatrist, and a candidate who can state them is well placed to answer the item. The first is the sequence itself. Because **the syndrome usually arrives before the malignancy has been detected**, the first clinician to see the patient is whichever clinician the presenting complaint routes them to. Where that complaint is behavioural, meaning altered mood, altered personality, psychosis, failing memory, reversed sleep, agitation or apathy, that clinician is a psychiatrist. The referral letter says nothing about cancer, because at that point there is nothing about cancer to say.

The second reason is that psychiatric formulations are generous, and their generosity is usually a strength. A subacute behavioural change in an adult can be absorbed comfortably into a first episode of psychosis, a late-onset depressive episode, a mood disorder cycling faster than usual, a dissociative presentation or a functional cognitive disorder. Each is internally coherent and each generates a plausible treatment plan. None of them obliges the clinician to ask what else could be producing this picture, so the question of occult malignancy is one the reasoning process will never raise unless the clinician raises it deliberately.

The third reason is what the examination is actually testing. The bank item asks when to suspect, and suspicion is a skill with a location in time: it belongs at first contact, alongside the initial formulation, and it loses value the longer it is deferred. A suspicion raised after several failed treatment trials is a lesser clinical act than the same suspicion raised at intake, because in the interval the syndrome has progressed while **the malignancy has continued to sit undetected**.

TRAP **Treating the difficulty as a laboratory problem.** Candidates picture the failure as a missed antibody result, so they write about panels and assays. The failure that actually happens is upstream and cheaper to fix: the question of an underlying tumour was never opened, so nothing was ever sent. An answer that begins with the test has skipped the step the item is examining.

14.3 Reading the presentation: tempo, company and fit

Suspicion is a discipline rather than a mood, and it should not be taught as vigilance in general. It is built from a small set of readings that a competent psychiatric history already produces, and the work is to read them deliberately rather than to hope they announce themselves.

The first reading is **tempo**. Psychiatric illnesses have characteristic time signatures, and a clinician who has taken many histories knows them implicitly. What matters here is the mismatch: a presentation that moves faster than the illness it resembles, that keeps moving while treatment is being delivered adequately, and whose trajectory is a slope rather than an episode.

Tempo costs nothing to record, requiring only a careful timeline with collateral corroboration, and it is the reading most often left vague in a psychiatric clerking.

The second reading is **company**, meaning what the psychiatric syndrome has arrived with. Psychiatric syndromes keep characteristic company, and neurological company is different company. When a behavioural presentation is accompanied by features belonging to the nervous system rather than to the mind, the formulation has acquired something it cannot metabolise, and that is informative in itself. The plate below sets out the clusters that carry this weight, and it is the right place to learn them, because they are the content of the source criteria rather than of this author's recollection.

The third reading is **fit**: whether this person fits the illness being diagnosed, and whether a general history has been taken at all. Psychiatric clerking has a known habit of collecting an exhaustive developmental history alongside a thin systemic one, so constitutional features, general medical risk and the ordinary questions an internist would ask are frequently missing. A history that never asked them cannot be reassuring, and the difference between a negative finding and an unasked question is one this chapter returns to repeatedly.

The plate that follows carries the working set: the onconeural antibody group, the tempo and sign clusters that raise suspicion, and the boundary against the autoimmune-encephalitis workup taught elsewhere. Read the plate as the operational list, and this prose as the reasoning that tells you when to reach for it.

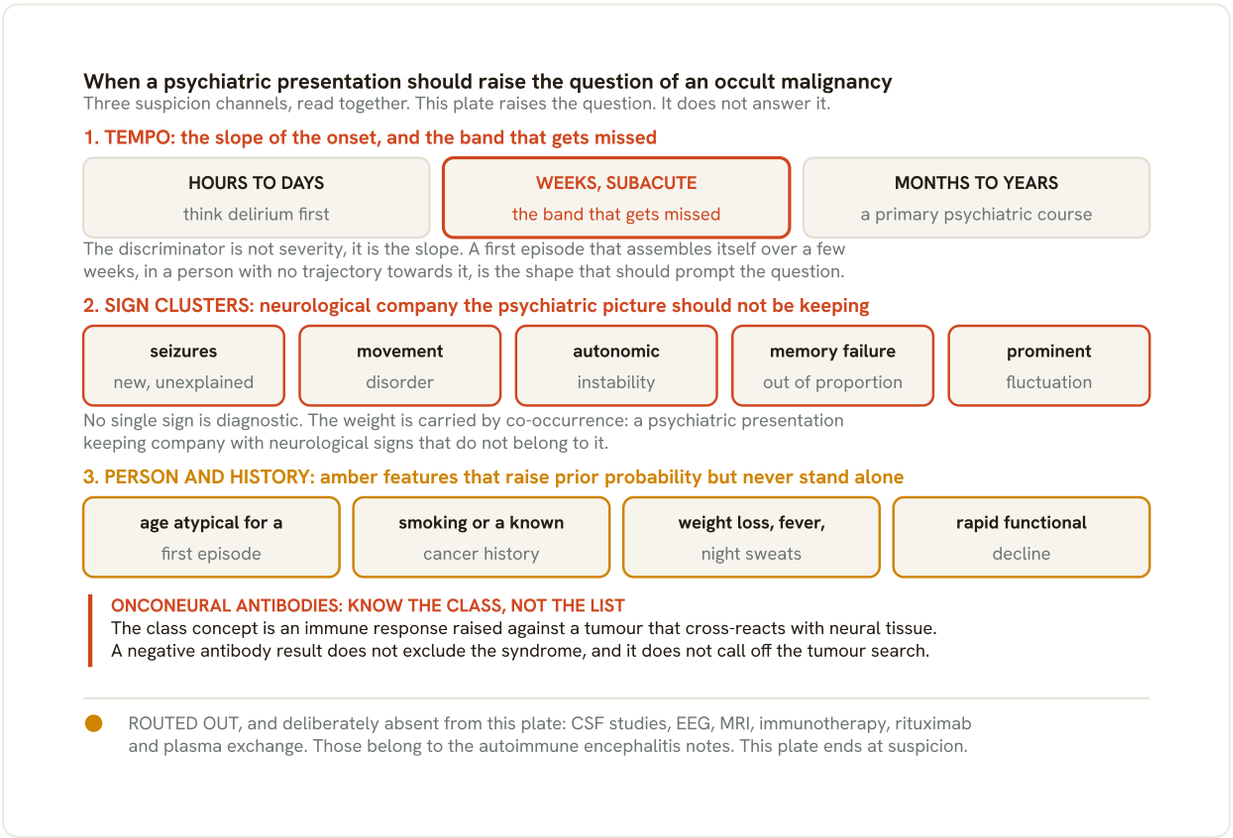


Fig 40.8 The three suspicion channels for an occult malignancy presenting psychiatrically. The plate is deliberately a suspicion instrument and stops there, because the workup and its immunotherapies belong to the autoimmune encephalitis notes and are cross-referenced out rather than duplicated. What it adds to the prose is the tempo rail, which puts the subacute band between two safer neighbours and shows why that band is the one that gets missed: a presentation over hours to days is read as delirium and gets investigated, a presentation over months to years is read as primary psychiatric and gets followed, and the few-week presentation falls between two reflexes. The onconeural antibody idea is carried as a class rather than a list because that is how it is examined and how it behaves, and because a negative result narrows nothing. From memory you should be able to state the tempo band that should worry you, name the sign clusters that a psychiatric picture should not be keeping company with, and say what a negative antibody result does and does not permit you to conclude.

14.4 The boundary: suspicion is yours, the workup is not

This section owns paraneoplastic suspicion. It does not own the autoimmune-encephalitis workup. Cerebrospinal fluid examination, electroencephalography, magnetic resonance imaging, first-line immunotherapy, rituximab, plasma exchange and intravenous immunoglobulin are taught in the Neuropsychiatry of systemic and neurological disease notes, and this text restates no parameter of any of them. The boundary is named explicitly because it is the one an author is most tempted to cross: the workup has protocols and sequences and feels like the substance of the topic, and it is the substance of a different topic with a different owner.

The split is clinical rather than merely editorial. The investigations and the immunotherapy form one coherent body of knowledge, learned and examined together. Suspicion is a separate skill, exercised by a different clinician at an earlier moment, and it is the skill this bank item names.

What the psychiatrist owns is worth stating positively, and it is short enough to hold as a list.

- The timeline, recorded precisely enough that tempo can be read off it later by a clinician who was not in the room.

- The collateral history, often the only source of that tempo.
- The general and systemic history that psychiatric clerking tends to abbreviate.
- The referral itself: one that says "for neurology opinion" transmits only the patient, while one that states the tempo, the company and the explicit reason a tumour is being considered transmits the reasoning.

THE QUESTION IN FRONT OF YOU	WHERE IT IS ANSWERED
Could this psychiatric presentation be driven by a tumour nobody has found yet?	This section.
What do the cerebrospinal fluid, the electroencephalogram and the scan show, and what immunotherapy follows?	Neurology-psychiatry interface notes, autoimmune encephalitis. Never taught here.
Is this delirium, and how is it assessed and managed?	Stroke, traumatic brain injury, delirium and focal brain syndromes notes.
Does this patient have capacity, and what is the statutory test?	Mental health legislation notes. Capacity is decision-specific and time-specific; the test is taught there.
How is the psychotropic adjusted for hepatic or renal impairment?	Psychopharmacology notes. This chapter adds no further version of a dose adjustment.
Is the disturbance caused by the cytotoxic agent, by interferon, or by the corticosteroid?	The sibling section on chemotherapy and interferon effects, and the routing section owning steroid-induced disturbance.
Is this neuropsychiatric lupus, neurosarcoidosis, central nervous system vasculitis or antiphospholipid syndrome?	The sibling sections owning each of those.
Is this depression, demoralisation or adjustment in someone already known to have cancer?	The routing section on cancer.

14.5 The screening imperative, and the guides it names

Once suspicion exists, the criteria are direct about what follows. **Cancer screening should be promptly undertaken when a paraneoplastic neurologic syndrome is suspected, and should be guided by the type of phenotype or antibody.** That single sentence is the operational core of this section, and it repays reading in parts.

The first operative word is **promptly**. It fixes a position in the sequence: the screen belongs at the point of suspicion rather than at the point of diagnostic certainty. The claim carries urgency and carries no interval, so this text supplies none, and a candidate who wants a period in days or weeks should take it from the criteria rather than from a note that would be inventing it.

The second operative phrase is **guided by the type of phenotype or antibody**. The instinct under pressure is to describe a broad sweep, imaging everything and hoping something appears. The criteria describe a search with a direction, where the clinical phenotype narrows the field of plausible primaries and the antibody, when available, narrows it further. An answer that reproduces the urgency while dropping the guidance has reproduced the easier part.

One inference is worth stating, because it is reasoning rather than recollection. If **the syndrome characteristically presents before the malignancy is detected**, the screen is being run in a body where the tumour has so far evaded detection. A first look that finds nothing has therefore not settled the question the way a first look normally settles one, and the posture that follows is continued attention rather than closure.

RULE **Prompt and guided, in that order and with equal weight. The criteria couple the timing of cancer screening to the suspicion of a paraneoplastic neurologic syndrome, and couple its direction to the phenotype or the antibody.** Reproduce the pairing. An answer carrying only urgency describes panic; an answer carrying only direction describes a screen that may arrive too late to matter.

Now the honest limit of this section, narrowed to what the ledger actually withholds. The antibody pairings are cited and taught in the subtopic below, so they have left the list of absences. What is still missing is the machinery around the screen: no row names which imaging modality corresponds to which phenotype, no row gives an interval for repeating a screen that found nothing, and no row states how long surveillance continues. Those three are what a candidate will want next, and they are owed to the source criteria rather than to this note.

14.6 Onconeural antibodies: the risk tier, the phenotype and the reported pairings

The term **onconeural antibody** names an antibody directed against an antigen shared by tumour tissue and nervous tissue. That definition is why the antibody is diagnostically useful at all: because the antigen belongs to each compartment, a serological finding points at a body region that has yet to be searched. Take the term seriously as vocabulary, because it carries the logic of the whole section in a single word.

The antibody earns its place here for the reason the criteria give: **the cancer screen is guided by the type of phenotype or antibody**. The antibody is a laboratory result while the syndrome is a clinical entity, so the result informs the search rather than constituting the diagnosis. And it sits alongside the phenotype in the claim as written rather than replacing it, so a clinician who has a phenotype and awaits a result already holds one of the guides the criteria name.

The criteria then sort the antibodies by the cancer risk each one carries, and that sorting is what makes the guidance operational. **The first group holds the high-risk antibodies, defined as those associated with an underlying cancer in more than 70 per cent of the patients who carry them.** A structural fact travels with the group: **most high-risk antibodies target intracellular antigens**. Put the two together and an inference follows. An antibody raised against an antigen sitting inside the cell is poorly placed to be doing the damage by binding it, so the finding is best used as a pointer towards a hidden tumour rather than as an account of mechanism.

One phenotype belongs to psychiatry more than to any other specialty. **Paraneoplastic limbic encephalitis is characterised by personality change, irritability, depression, seizures, memory loss and sometimes dementia.** Read that list as a psychiatric referral rather than a neurological one. Personality change, irritability and depression are what a psychiatric service is

built to receive; the seizure and the memory failure are what eventually move the formulation, often long after a psychiatric label has been applied and acted upon.

The reported pairings follow, with the evidence grade placed second so that it is read before the pairing it qualifies, and with the two rows that rest on more than one patient placed first. Read the second column before the third. A single reported case establishes that a pairing has been observed in a human being, and nothing beyond it: no frequency, no claim that this tumour is the one the antibody usually points to, and no entitlement to the word "classical".

ANTIBODY	WHAT THE CITED ROW IS	SYNDROME REPORTED WITH IT	TUMOUR REPORTED WITH IT
CRMP5	**The commonest associated diagnoses in a 24-patient cohort.**	**Peripheral neuropathy and cerebellar ataxia**	**Lung cancer**
Anti-Ma2	**Described as a specific marker for that pairing.**	**Limbic and brainstem encephalitis**	**Testicular germ-cell tumour, mainly in adult males**
Anti-Hu	**A single reported case.**	**Adie's pupil with sensorimotor polyneuropathy**	**Primary mediastinal small-cell carcinoma**
Anti-Yo	**A single reported case, arising years after the tumour was resected.**	**Cerebellar degeneration**	**Ovarian cancer**
Anti-Ri	**A single reported case.**	**Opsoclonus-myoclonus**	**Breast carcinoma, in a male patient**
Anti-amphiphysin	**A single reported case.**	**Stiff-limb syndrome**	**Invasive ductal carcinoma of the breast**

The second column is where marks are won and lost, which is why it now sits where the eye reaches it first. Four of these six rows rest on one patient each, and several of those reports were published precisely because they were unusual: an opsoclonus-myoclonus syndrome in a man, a cerebellar syndrome declaring itself years after the ovarian tumour was resected, a mediastinal rather than a pulmonary small-cell primary. A report selected for its rarity is the worst available basis for a statement about what is typical, and an answer that flattens this table into a list of classical associations has manufactured a confidence its own sources never carried.

The criteria also grade how much confidence the label deserves. **The PNS-Care score is a composite model for assigning diagnostic likelihood in a patient with a suspected paraneoplastic neurological syndrome**, and it answers in bands: **definite at 8 or above, probable at 6 to 7, possible at 4 to 5, and a score of 3 or less classed as something other than a paraneoplastic syndrome**. Learn the bands, and learn that a graded answer exists. The items the score weighs are absent from this fragment's ledger, so they are left to the criteria.

Syndrome, then tumour	Prompt and guided	Pointer, not diagnosis	Observed, not typical
THE SEQUENCE THAT DEFINES THE TERM	HOW THE CANCER SCREEN IS RUN	WHAT AN ONCONEURAL ANTIBODY IS FOR	WHAT A SINGLE-CASE PAIRING ROW ESTABLISHES

The four tiles carry the spine of the section: a sequence, a manner of screening, a use for the laboratory result, and a limit on what the pairing table licenses.

IN INDIA

The screening imperative assumes an onconeural antibody panel can be requested, paid for and reported within a useful period, and in much of Indian practice each assumption deserves checking before the plan is written. Panels are sent out to referral laboratories from most district and many teaching hospital settings, the cost is usually met out of pocket, and turnaround runs to the time a courier and a batch run take. The practical consequence is that **the phenotype is frequently the guide available first, with the antibody arriving later to refine the search**, which makes a carefully documented clinical picture more valuable here than where serology returns quickly. The Indian record on availability, pricing and reporting intervals for these panels was never located within this fragment's ledger, so it is described qualitatively rather than quantified.

GOOD TO REMEMBER

On the ward this reduces to a single habit. When a psychiatric presentation moves at the wrong speed, keeps company with neurological features, and sits in a person whose general history was never properly taken, write the timeline down and say the word "paraneoplastic" out loud in the referral. **The syndrome characteristically arrives before the malignancy is found**, so the psychiatrist is often the clinician best positioned to raise the question at the moment it is most useful. What follows, the fluid, the imaging, the electrophysiology and the immunotherapy, belongs to colleagues and to another section. Raising the question well is the part that is yours.

The syndrome almost always arrives before the malignancy is found, so the psychiatrist's part is suspicion at first contact: read tempo, company and fit, write the timeline down, say the word in the referral, and let the cancer screen be prompt and guided while the workup belongs to somebody else.

Autoimmune and inflammatory systemic disease, and burn injury

15 · Neuropsychiatric systemic lupus erythematosus: manifestations, diagnosis and management

The examinable difficulty in neuropsychiatric lupus is not the list of syndromes, it is deciding whether the syndrome in front of you belongs to the disease at all. Lupus reaches the nervous system often enough to produce almost any neurological or psychiatric presentation, and the same patient can also be septic, uraemic, hypertensive, taking drugs with central effects, and

frightened by an illness that has already taken things from them. This section installs the topic in the order the decisions are made: what the term covers, how the manifestations reach a psychiatrist, how attribution is reasoned when no test settles it, and what treatment looks like afterwards.

Several things this presentation touches are owned elsewhere and are named here rather than rebuilt. Delirium assessment and management, including the bedside instruments and the drug choice for an agitated delirious patient, belong to the stroke, traumatic brain injury, delirium and focal brain syndromes notes. Organ-adjusted prescribing belongs to the psychopharmacology notes, and the workup of autoimmune encephalitis to the neurology and psychiatry interface notes. Corticosteroid-induced psychiatric disturbance is taught by the routing section of this chapter that owns it and by the endocrine and metabolic notes, which matters here because the drug class used to treat this disease is itself a recognised cause of psychiatric disturbance. The antiphospholipid syndrome is owned by the sibling section that carries it alongside neurosarcoidosis and central nervous system vasculitis.

15.1 What the term is asked to cover, and why its breadth is the problem

Neuropsychiatric systemic lupus erythematosus, conventionally shortened to **NPSLE**, is a collective label for the neurological and psychiatric manifestations occurring in a person who has lupus. It is a container, not a diagnosis of the kind a psychiatrist is used to making. Standard usage sorts its contents into central nervous system events, peripheral nervous system events and psychiatric syndromes, and those groups share almost nothing at the level of mechanism, severity, investigation or treatment.

That breadth is the source of every difficulty that follows. A label covering both a mild headache and a catatonic psychosis cannot have one test, one mechanism or one treatment, and any sentence beginning "NPSLE is treated with" is suspect before its predicate arrives. The useful move is to handle the label as a question in parts: which syndrome, by which mechanism, in a patient whose disease is at what level of activity, with which alternative explanations excluded.

RULE **Name the syndrome, then the mechanism, then the attribution, then the treatment.** These are separate determinations and each can be right while the next is wrong. A candidate who says "this is neuropsychiatric lupus" has named a container. One who says "this is an acute confusional state, in a patient with active disease, in whom infection and metabolic causes have been excluded, treated as an inflammatory manifestation" can be examined on any of the four.

15.2 How the manifestations reach a psychiatrist

Grouping the manifestations by the route through which they arrive is more useful than grouping them by anatomy, because the route determines what you are being asked to do.

- **Diffuse presentations**, which psychiatry is most often called about: an acute confusional state, mood disturbance of either polarity, psychosis, cognitive impairment across attention, memory and executive function, and anxiety. Here the referral question is usually not "is this lupus" but "is this psychiatric".

- **Focal central events**: seizures, stroke-like episodes, movement disorder, myelopathy, aseptic meningitis, cranial neuropathy and a demyelinating syndrome. These reach neurology or medicine first, and psychiatry is called for the behavioural consequences, for capacity, or because the presentation was read as functional.
- **Peripheral events**: polyneuropathy, mononeuropathy, plexopathy, autonomic disturbance and a neuromuscular junction disorder. These rarely present to psychiatry directly, and their relevance is that pain and disability from them drive mood.

The classes are nomenclature rather than mechanism: they say where the disturbance sits and nothing about why. This chapter also holds no row giving the frequency of any manifestation in any population, so nothing here is called the commonest and no ordering by frequency is printed. Frequency statements circulate widely in this topic and disagree with one another, for the reason the next subtopic makes clear.

15.3 Attribution, which is the actual core of the topic

Everything listed above can be produced by lupus, and every one of them can be produced by something else in a person who happens to have lupus. Diagnosis is therefore **attribution**: a reasoned judgement that this syndrome, in this patient, at this time, belongs to the disease. It is reasoned rather than measured, because **there is no diagnostic gold standard by which the neuropsychiatric syndromes of lupus can be diagnosed**. That fact is the axis of the topic and belongs early in any answer.

Three consequences follow, each worth a sentence. Published frequency estimates disagree, because without a gold standard every study defines its own numerator: a series counting headache and mild mood symptoms and a series counting only severe syndromes are not measuring the same thing, and their percentages are not comparable even when both are correctly reported. No investigation is confirmatory, so every result is evidence to be weighed rather than an answer. And attribution is revisable, since a judgement made on admission may be overturned by the culture result, by the course, or by the response to treatment.

The reasoning runs on a small number of axes worth naming. The first is the temporal relation to disease activity elsewhere: a syndrome appearing while the disease is active in other organs is easier to attribute than one appearing in quiescent disease. The second is the presence or absence of a competing explanation, which the next subtopic takes up because it does the most work. The third is the character of the syndrome, since a course, phenomenology or examination finding atypical for a primary psychiatric disorder points towards a systemic cause. The fourth is response to treatment, and it is retrospective, which is why the attribution has to be made and acted on before the evidence for it is complete.

TRAP **Treating a positive antibody as an attribution.** Serology establishes that the patient has the disease. It does not establish that the syndrome in front of you was produced by it, and a patient with unambiguous lupus can have an unrelated psychiatric illness, an infection or a drug effect. Skipping the attribution step is also how a treatable cause is missed while immunosuppression is escalated.

15.4 What has to be excluded before the disease is credited

Because attribution rests so heavily on the absence of a competing explanation, the exclusion list is the substance of the assessment rather than preliminary housekeeping. Infection comes first and stays first in a patient immunosuppressed by both the disease and its treatment, and it includes meningitis, encephalitis, and the encephalopathy of systemic sepsis without central infection. Metabolic causes follow: uraemia, sodium and glucose disturbance, hepatic and thyroid disturbance. Hypertensive encephalopathy belongs in anyone with renal disease and a new confusional or seizure presentation.

Drug causes are the largest single trap here, because the treatment of the disease and the cause of the psychiatric syndrome can be the same agent. That construct is taught by the routing section of this chapter that owns it; what belongs here is the boundary. A new psychiatric syndrome beginning after a change in immunosuppressive treatment, in a patient whose disease activity elsewhere has improved, pulls in the opposite direction from an attribution to the disease. The distinction changes treatment completely, since one path escalates immunosuppression and the other reduces it.

Primary psychiatric illness stays on the list throughout, supported by a psychiatric history predating the lupus, a family history, and a typical course. So does the ordinary burden of a relapsing systemic illness with a visible skin component and a treatment with visible adverse effects: demoralisation and adjustment disorder in that setting are a person responding to their circumstances rather than the disease reaching the brain, and are treated as such.

Finally, lupus is one member of a family of systemic diseases that reach the nervous system, and the discriminating features are what separate them at the bedside. The plate accompanying this section sets lupus beside the other systemic autoimmune presentations this chapter teaches, neurosarcoidosis, central nervous system vasculitis, the antiphospholipid syndrome, coeliac disease and inflammatory bowel disease, each with its discriminating feature. Each is taught in full by the sibling section that owns it and is named here only to place lupus in its family.

Six systemic autoimmune conditions on one comparator, because the task is telling them apart
One row each. The middle column is the only one that separates them from each other.

CONDITION	THE DISCRIMINATOR, in one feature	WHAT BRINGS IT TO PSYCHIATRIC NOTICE
Neuropsychiatric lupus	multisystem: skin, joints, kidney and blood counts, all in the one patient	psychosis or a mood episode arising inside an already declared illness
Neurosarcoidosis	cranial neuropathy, with basal meningeal or pituitary involvement	apathy and personality change, with hypersomnia and endocrine mood shift
Primary CNS vasculitis	confined to the CNS, with headache and stepwise deficits on a bland screen	no psychiatric presentation is claimed here: a declared gap, see the text
Antiphospholipid syndrome	thrombosis and pregnancy loss are the carrying features, often recurrent	cognitive complaints in a young adult with a thrombotic or obstetric history
Coeliac disease	gluten-driven enteropathy, and the deficiency load does much of the work	depression and apathy that persist after the mood treatment is optimised
Inflammatory bowel disease	relapsing bowel disease, and the burden tracks disease activity	depression and anxiety that flare with the bowel, plus steroid-driven change

READ DOWN THE MIDDLE COLUMN: that is the only column that closes the differential.
Low mood and cognitive change run across these rows, so that column raises the question but cannot answer it.

ROUTED OUT: the serological and imaging workup for each of these six, and their immunosuppressive treatment, are not taught on this plate. What is taught here is telling the six apart.

Fig 40.9 Six systemic autoimmune conditions laid on a single comparator so that they can be told apart rather than revised one at a time. The plate exists because the right-hand column is nearly identical down all six rows, and that is the teaching point: depression, anxiety and cognitive change are the common currency of systemic inflammation, so the presentation that brings the patient to psychiatric attention almost never identifies which condition it is. The middle column carries the whole differential, and each entry is chosen to be true of that row and false of the other five. One right-hand cell is deliberately empty of a clinical picture: no source explicitly linking primary central nervous system vasculitis to a psychiatric manifestation was located for this chapter, so the gap is printed rather than filled from a neighbouring condition, and the accompanying text says what to do with it. Note the two structural families the table exposes, that the first four reach the brain directly through inflammation, vasculopathy or thrombosis, while the last two reach it indirectly through nutrition, disease activity and treatment exposure. From memory you should be able to give the single discriminating feature for each of the six, and say why the psychiatric column can raise the question but never settle it.

15.5 Investigation, and what each result can and cannot settle

The investigation logic follows from the absence of a gold standard: every test is chosen for what it excludes or supports, and none to confirm. Serology and disease activity markers establish that lupus is present and give some sense of activity elsewhere, which feeds the temporal axis and nothing more. Structural imaging looks for a focal lesion, for ischaemic change and for an alternative structural explanation, and a normal scan is entirely compatible with a severe diffuse syndrome, which is the most misread result in this topic.

Lumbar puncture here is done principally to exclude infection, and its findings in the disease itself are non-specific. Electroencephalography supports or refutes a seizure explanation for an episodic disturbance and characterises an encephalopathy without attributing it. Neuropsychological assessment characterises a cognitive deficit, which is valuable for

rehabilitation and for tracking change, and does not attribute it. Where autoimmune encephalitis is genuinely in question, that workup belongs to the neurology and psychiatry interface notes.

This chapter holds no row carrying an antibody to syndrome association, an imaging finding to syndrome association, or the performance of any investigation in this population, so none is asserted. That is a real gap rather than a stylistic omission, and the antibody associations should be read from a primary source and held separately from what is taught here.

RULE **No investigation closes the question, because no diagnostic gold standard exists for these syndromes.** Order tests to exclude the competing explanations and to characterise the syndrome, and expect to make the attribution on the balance of what you have. An answer promising a confirmatory test is wrong on the central fact of the topic.

15.6 Treatment, and the recommendation this section carries

Once a syndrome has been attributed to the disease, treatment has three limbs, chosen by mechanism rather than by phenomenology. The **inflammatory limb** is immunosuppression, applied when the syndrome is judged to arise from active disease. The **thrombotic limb** is antithrombotic treatment, applied when the mechanism is judged to be vascular occlusion; the antiphospholipid syndrome that most often drives it is owned by the sibling section, so this section names the limb and stops. The **symptomatic limb** is psychiatric and neurological treatment of the syndrome itself, running alongside either of the others rather than after them.

The guideline recommendation this section carries is precise about its population and about its comparison, and both matter more than the drug names. In adult patients with lupus-related severe, acute neuropsychiatric manifestations, the guidance is to **use GCs plus CYC over GCs alone or GCs plus RTX**. The abbreviations are the guideline's own and stand for glucocorticoids, cyclophosphamide and rituximab.

Read the population before using the recommendation, because it carries restrictions a candidate quoting it loosely will drop. It concerns adults, so it says nothing about children. It concerns severe manifestations, so it says nothing about the many mild ones this container also holds. It concerns acute manifestations, so it says nothing about chronic or slowly progressive syndromes, and a chronic cognitive syndrome is precisely what psychiatry is most likely to be asked about. Quoted as though it applied to everyone with this label, it becomes an error of scope.

What this section does not carry is equally worth stating: no dose, no route, no duration, no schedule, no maintenance strategy, no comparison for the mild or chronic presentation, and no statement about what to do when the first choice fails. Organ-adjusted prescribing for renal involvement belongs to the psychopharmacology notes, and remembered regimens in this topic diverge sharply between sources.

15.7 The psychiatric limb, and the boundaries around it

Psychiatric treatment of an attributed syndrome runs alongside the immunosuppression and is not deferred until it works. An acute confusional state is managed as an acute confusional state, by the pathway taught in the delirium notes, while the cause is treated. A psychotic or affective

syndrome is treated symptomatically with two standing cautions in view: renal and haematological involvement constrain the choice of agent, quantified in the psychopharmacology notes rather than here, and the possibility that treatment of the disease is contributing keeps the drug review open.

Two further boundaries belong in an answer. Capacity is assessed as it is anywhere else, and the only things this section says about it are that it is decision-specific and time-specific, which matter unusually often in a fluctuating encephalopathy; the doctrine and the statutory test belong to the mental health legislation notes. Adherence deserves explicit attention, because the psychiatric syndrome sits on top of a demanding medical regimen, and a cognitive syndrome never named will be recorded as non-compliance.

The rehabilitative and psychological work is most often omitted and has the longest reach. A cognitive deficit that has been characterised can be worked with, at home and at work, whether or not it is reversible, and the demoralisation that follows a burdensome illness responds to being named and addressed rather than to being reclassified as a manifestation of the disease.

15.8 Answering the question

The question is almost always some form of "a patient with known lupus presents with a psychiatric syndrome, discuss". Each row below is a decision rather than a heading, so an answer that walks the rows is reasoning aloud rather than reciting.

DECISION	WHAT MOVES IT	WHAT CANNOT SETTLE IT	WHERE THE DETAIL IS TAUGHT
Which syndrome is this	Phenomenology, course, examination, cognitive profile	Serology, which speaks to the disease and not the syndrome	Here
Is it attributable to the disease	Timing against disease activity, absence of a competing explanation, atypical features	Any single investigation, since no confirmatory test exists	Here
Is it infection, metabolic or hypertensive	Cultures, biochemistry, blood pressure, imaging, lumbar puncture	The presence of lupus, which raises infection risk	Here, and the delirium notes
Is it the treatment rather than the disease	Onset after a treatment change, disease quiescent elsewhere, the drug review	Severity, which is uninformative about cause	The routing section owning the corticosteroid construct
Inflammatory or thrombotic mechanism	Focal against diffuse presentation, vascular imaging, antibody profile	Response to a treatment not yet given	Here for the limbs, the sibling section for the antiphospholipid syndrome
What is given, and how is it prescribed	Severity and acuity of the manifestation, and the age group	A remembered dose or regimen	Here for the recommendation, psychopharmacology for prescribing

IN INDIA

No Indian row was located for this item, and the shortfall is declared rather than filled: no Indian cohort, no Indian frequency estimate, no Indian figure for the cost or availability of either named agent, and no Indian guideline statement on this manifestation. The recommendation this section does carry comes from a Latin American clinical practice guideline, so its transferability to an Indian public hospital is a judgement rather than a given, and a candidate who names that is on stronger ground than one who quotes it as though it were local. What is safe to say in an Indian answer is structural: infection carries the most weight of the exclusions where the background burden of chronic infection is high, immunosuppression is decided jointly with rheumatology, and the rehabilitative limb is the one most likely to be dropped where follow-up is difficult.

PITFALL

Answering as though attribution were a formality on the way to treatment. An answer spent on immunosuppression has skipped the contested part, and also the part that protects the patient, since the competing explanations are commoner and more immediately reversible than the manifestation being reached for. Spend the answer where the difficulty is.

GOOD TO REMEMBER

Called to a ward for a patient with lupus and a new psychiatric presentation, give the medical team your view on whether the disease has reached the brain as a judgement rather than a finding, and record what would change it. An attribution offered in that form can be revised without anybody losing face; one asserted as a finding is very hard to withdraw.

MNEMONIC

Syndrome, Sepsis, Salts, Script, Then Attribute

Characterise the syndrome, exclude sepsis and infection generally, check the salts and the rest of the metabolic screen, read the script for drugs that could be doing it, and attribute to the disease last. Everything before the last step is faster, commoner and more reversible than what comes after it.

No gold standard	Attribution	GCs plus CYC
FOR DIAGNOSING THE NEUROPSYCHIATRIC SYNDROMES OF LUPUS	WHAT DIAGNOSIS ACTUALLY CONSISTS OF HERE	PREFERRED OVER GCS ALONE OR GCS PLUS RTX, ADULTS, SEVERE ACUTE MANIFESTATIONS

There is no test that confirms neuropsychiatric lupus, so diagnosis is an attribution built from timing, exclusion and the character of the syndrome, and the exclusions that matter most are infection, metabolic disturbance and the treatment itself. Treatment then splits by mechanism into an inflammatory limb, a thrombotic limb and a symptomatic limb, and the one recommendation this section carries applies to adults with severe, acute manifestations and to nobody else.

16 · Neurosarcoidosis, CNS vasculitis and antiphospholipid syndrome with psychiatric features

Neurosarcoidosis, primary vasculitis of the central nervous system and antiphospholipid syndrome are filed together as rare systemic causes of psychiatric presentation, and that filing is very nearly the only thing they share. Each arrives in psychiatry by a different route, each is settled by a different kind of evidence, and for one of them the psychiatric literature this text could locate is thin enough that saying so is the correct answer rather than an embarrassment. This section installs what is actually sourced for neurosarcoidosis, what is actually sourced for antiphospholipid syndrome, and a worked method for answering on primary central nervous system vasculitis without inventing a phenotype for it. The rare causes are examinable for exactly this reason: they are where a candidate's discipline about evidence becomes visible inside a single paragraph.

16.1 Three names, three different examination problems

The grouping is administrative. A syllabus has to put the uncommon systemic causes somewhere, and the shared shelf tempts a tired candidate into one undifferentiated answer about inflammation reaching the brain. That answer earns very little. The examiner assumes you accept that systemic autoimmune disease can affect the mind, and is asking something else: whether you can tell these conditions apart with a patient in front of you, and whether you know which parts of the answer are yours to give.

It helps to sort them on a single axis: where the discriminating information sits. For **neurosarcoidosis**, the sourced material puts discriminating information inside the mental state itself, because the described presenting picture assembles features that a purely functional psychosis would not usually assemble in the same person at the same time. For **antiphospholipid syndrome**, the sourced material does the opposite: the reported psychiatric syndromes are the ordinary syndromes of general adult psychiatry, so the mental state is precisely where the discrimination is not. For **primary central nervous system vasculitis**, the discriminating information sits in a specialty literature that this text could not quote, and that changes the shape of an honest answer without making the condition irrelevant.

Those are different problems and they take different answers. The first is a recognition problem: can you notice, in a mental state you are already examining, a combination that sends you outside psychiatry. The second is an attribution problem: with a systemic diagnosis on the file and an ordinary psychiatric syndrome in front of you, how much may you claim about the relation between them. The third is an evidence problem: what do you say when the psychiatric claim you would like to make is one your sources will not carry. A candidate who answers all three with one paragraph about immune-mediated brain injury has told the examiner that the differences are invisible to them. The generic immune-to-brain framework belongs elsewhere in this volume and is deliberately not rebuilt here.

16.2 Neurosarcoidosis: a picture that can arrive before the diagnosis

The sourced statement for neurosarcoidosis is a statement about presentation. Reported neuropsychiatric manifestations include **confusion, dysfunction in orientation or memory, hallucinations, delusions and catatonia**, and these are described as manifestations that can be the presenting features of the condition. Read the load-bearing word carefully. Presenting means the psychiatric picture can be the first thing anyone sees, which in turn means the referral may reach a psychiatrist while the systemic diagnosis is still unmade, unsuspected and absent from the file you are handed.

That is the examinable point, and it is about sequence rather than rarity. Were the systemic diagnosis always established first, this condition would be a footnote in liaison practice, because the medical team would carry it. Because the neuropsychiatric picture can come first, the recognition burden lands on psychiatry, during an ordinary assessment of what looks like a first episode of psychosis.

The discriminating feature is the combination rather than any single item. The sourced list runs across layers that psychiatric nosology usually keeps apart: a confusional and cognitive layer (**confusion, dysfunction in orientation or memory**), a psychotic layer (hallucinations and delusions), and a motor and behavioural layer in the form of catatonia. A patient in whom orientation and memory are disturbed, psychotic phenomena are present, and catatonic signs are observable has a picture crossing those layers at once. Crossing them is the signal, and the response is a decision to widen the question beyond the psychiatry room rather than a diagnosis made inside it.

Be equally precise about what the source does not license. The material behind that list is a report of a single presentation, and its framing describes these manifestations as rare. A statement of that kind establishes possibility. It does not establish frequency, it does not rank the features against one another, and it supplies no proportion of patients in whom any of them occurs. So there is no percentage in this section, and a candidate who supplies one has supplied it from memory. The marks here are for recognising the pattern and routing the patient, never for a prevalence figure.

RULE **A presenting feature is not a common feature.** The sourced claim about neurosarcoidosis licenses recognition: you keep the condition available when a picture crosses cognitive, psychotic and catatonic layers together. It licenses no statement about how often that happens, and none that this presentation is typical of the disease. State the possibility, then say what you would do about it.

16.3 What psychiatry contributes when neurosarcoidosis is the question

Once the question is open, the psychiatrist's contribution is a description good enough for someone else to act on. Record phenomenology in the patient's own words rather than in summary labels, because the summary label is what a later reader will doubt. Record catatonic phenomena as observed signs with the manoeuvres that elicited them, so that a neurologist reading your entry can see what you saw. Date the cognitive change against the psychiatric history: a disturbance of orientation and memory that arrived with the psychotic phenomena

reads very differently from one that has been present for years, and only the psychiatric assessment is positioned to establish that ordering.

The rest of the pathway belongs to other owners, and naming them accurately is part of a complete answer. Imaging and cerebrospinal fluid evaluation of a suspected inflammatory brain process, with immunotherapy decisions, sits with the neurology and psychiatry interface material. Assessment and management of delirium as a syndrome, including its bedside instruments, sits with the delirium and focal brain syndromes material. Corticosteroid-related psychiatric disturbance, relevant once systemic treatment begins, is owned by its own routing section and by the endocrine and metabolic material. Adjustment of psychotropic prescribing for organ impairment sits with the psychopharmacology material. Capacity, if it arises, is decision-specific and time-specific, and the statutory test belongs to the mental health legislation material.

Routing is not evasion. An answer that names the right owner for each limb of the work shows the examiner how you manage a shared patient, and in particular that you know how each team assumes the other has already done the thinking.

TRAP **Recording catatonia as a diagnosis rather than as a set of signs.** Once catatonia is entered as a conclusion, the file behaves as though the question has been answered, and the systemic question quietly stops being asked. Candidates fall into this because catatonia feels like an endpoint: it is a named syndrome, it has a management pathway, and naming it produces the relief of having decided something. Written as observed signs instead, it stays what it is, which is a finding that still needs an explanation.

16.4 Primary central nervous system vasculitis: a declared gap, and how to answer it

Here this text has to say something uncomfortable and say it precisely. What was sought was a self-contained quotation from a primary study, a guideline or a systematic review that explicitly links primary central nervous system vasculitis to a psychiatric manifestation. No such quotation was located for this fragment. That is the gap, stated with its exact boundaries, and the boundaries matter: the search was for an explicit link to a psychiatric manifestation, quotable on its own, in a source of that standing.

Note what the gap is not. It is not a finding that the condition has no psychiatric presentation, and converting it into that claim would be a worse error than the one it replaces, because a categorical negative asserts far more than a failed search can support. Absence of a located quotation is a fact about this text's sources. It is not a fact about the disease.

Say also whose claim it would be. A statement linking this condition to a specific psychiatric syndrome would rest on the authors of peer-reviewed clinical studies of the condition and on the neurology and rheumatology specialty literature, where cohorts of these patients are described. Psychiatry cannot assume it by analogy with the conditions on either side of it in this section. The specialty that follows these patients owns the phenotype, and the honest position is to defer to it by name rather than fill the space.

Clinically the gap changes little about what you do and a great deal about what you write. The condition stays available as an inflammatory cerebrovascular possibility in a patient whose

presentation has crossed into neurological territory, and its evaluation sits with neurology as the neurosarcoidosis workup does. What changes is that you refrain from attaching a named psychiatric syndrome to it in a case note, a referral letter or an examination answer, because you would be attributing on authority you lack. You describe the presentation, record why the possibility is open, and let the specialty that can settle it settle it.

This is worth practising as a sentence, because candidates who have never said it aloud tend to bluff instead. A usable form runs: the psychiatric phenotype of this condition is not something I can state from a specific source, the neurology and rheumatology literature holds it, and my contribution is the described presentation and the referral. That answer is short, it is true, and it displays the judgement that separates a safe registrar from an unsafe one. An invented list of psychiatric features, delivered fluently, is the failure mode this discipline exists to prevent.

TRAP **Filling the vasculitis gap with borrowed lupus material.** Neuropsychiatric systemic lupus erythematosus has a rich catalogue of manifestations, it is taught in its own section of this chapter, and it sits close enough in a candidate's mind that its phenotype slides across to fill an empty space. It is a borrowed answer about a different disease, and an examiner who knows the vasculitis literature is thin will hear the borrowing at once. Name the lupus material as the neighbouring topic it is, and leave its content where it belongs.

16.5 Antiphospholipid syndrome: a psychiatric range that reads like general psychiatry

The sourced position for antiphospholipid syndrome is that a range of psychiatric disorders has been reported in these patients, and the reported range includes **psychosis, mania, depression, bipolar disorders and schizophrenia**. Read that list as a whole before reading any item in it. Taken together, **psychosis, mania, depression, bipolar disorders and schizophrenia** are close to a roster of general adult psychiatry. That is the discriminating feature of this condition, and it is a negative one in a useful sense: there is no distinctive mental state signature here to look for.

The consequence follows directly. If the reported psychiatric range is indistinguishable by name from primary psychiatric illness, the mental state examination cannot be the instrument that raises the possibility, and the discrimination has to come from elsewhere in the record. The syndrome is named for an antibody, which tells you that serology and the general medical history are where the question is settled, and that the settling belongs to the rheumatology and haematology literature. This text carries no sourced criteria for the condition, so it prints none, and a candidate who prints a criteria set from memory is doing the thing this method exists to stop.

Be careful with the verb as well. The sourced claim is that these disorders have been reported in patients with the condition. Reported in is an observation of co-occurrence in a described population, rather than a causal claim about the antibody producing the syndrome, and it carries no proportion. So a psychiatrist assessing a patient who already holds this systemic diagnosis has reason to keep the association in view while treating aetiology as unsettled. The material behind the claim is a narrative review drawing on a long publication window: a reasonable authority for the statement that a range has been reported, and a poor one for anything quantitative. There is no figure in this section either.

What this leaves you with is a clinical posture rather than an algorithm. In a patient with the systemic diagnosis on file who develops a new psychiatric syndrome, each of the easy errors is a real one. Treating the psychiatric presentation as unrelated by default discards information the record already holds. Attributing it to the systemic condition by default converts a reported association into a diagnosis and ends an assessment that had barely begun. The defensible position lies between them: keep the systemic diagnosis as a live question in the formulation, put that question in writing, and manage the psychiatric presentation on its own terms while it stays open.

RULE **Reported in is not caused by.** This distinction carries more marks across the chapter than any single disease fact, because almost every systemic association in liaison psychiatry is reported at the level of co-occurrence. Say what your source says, with the verb your source uses, then say separately what you would do clinically. An answer that keeps the moves apart reads as trained. An answer that fuses them reads as memorised.

16.6 Telling them apart in one pass

The plate for this section places the disease-specific presentations beside one another, including the neighbouring topics this section deliberately leaves to their owners: neuropsychiatric lupus, and coeliac disease and inflammatory bowel disease with the gut-brain axis. Use it as a recall grid in revision, working across the row for each condition and asking what makes that condition itself rather than the one above it. The table below does something complementary: it records where each condition in this section stands in evidence, which governs how firmly you may speak.

CONDITION	WHAT THIS TEXT CAN STATE	WHERE THE REMAINDER IS HELD
Neurosarcoidosis	That **confusion, dysfunction in orientation or memory, hallucinations, delusions and catatonia** can be presenting features, so psychiatry may see the patient first	Frequency, imaging and cerebrospinal fluid evaluation, and immunotherapy: neurology and the interface material
Primary central nervous system vasculitis	A declared gap: no self-contained quotation explicitly linking it to a psychiatric manifestation was located for this fragment	The psychiatric phenotype, if described: authors of peer-reviewed clinical studies and the neurology and rheumatology literature
Antiphospholipid syndrome	That a range of psychiatric disorders has been reported, including **psychosis, mania, depression, bipolar disorders and schizophrenia**	Diagnostic criteria, serology and proportions: the rheumatology and haematology literature

16.7 Answering on the rare causes without over-claiming

A complete answer here has a recognisable shape, worth rehearsing because the content is thin enough that structure carries the marks. Five moves carry it.

- Name the three conditions separately, and refuse the single undifferentiated paragraph about inflammation reaching the brain that the shared shelf invites.
- Neurosarcoidosis: give the discriminating feature as a presenting picture crossing cognitive, psychotic and catatonic layers together, and say that the response is to widen the question rather than to make the diagnosis in the psychiatry room.
- Antiphospholipid syndrome: give the discriminating feature as a reported psychiatric range that reads like general adult psychiatry, so the discrimination lives outside the mental state, in the record and the serology.
- Primary central nervous system vasculitis: give the discriminating feature as a psychiatric literature this text could not quote, held by the specialties that follow these patients, and name it as theirs rather than filling the space.
- Mark explicitly which parts of your answer rest on a source and which rest on clinical reasoning, then route the workup to its owner.

Recognition	**Attribution**	**Evidence**	**Reported, not caused**
THE NEUROSARCOIDOSIS PROBLEM	THE ANTIPHOSPHOLIPID PROBLEM	THE VASCULITIS PROBLEM	THE VERB THAT CARRIES THE MARKS

The first three tiles are the three problems this section separates, one to each condition, and a candidate who can say which of the three a question is asking has already done most of the discriminating. The fourth is the verb discipline that runs underneath all three.

The shape survives contact with a question you were not expecting, which is its real value. The neighbouring topics in this chapter, paraneoplastic neuropsychiatric syndromes, neuropsychiatric lupus, immunosuppressant neurotoxicity after transplantation, and the gut-brain topics, are owned elsewhere in this volume and each carries its own sourced material. Name them as neighbours when a question invites the comparison, and resist teaching them inside an answer about these conditions. Examiners read borrowed content as padding, and padding is expensive in a viva.

IN INDIA

No Indian series or Indian guidance on the psychiatric presentation of these conditions was located for this fragment, and that is a limit of this text rather than a claim about Indian practice or the Indian literature. What travels regardless of setting is the referral, and it is worth writing well. Specialist assessment may be separated in time from the presentation you witnessed, so the phenomenology, the catatonic signs with their eliciting manoeuvres, and the dating of cognitive change against the psychiatric history are things only your entry preserves. Where rheumatology, neurology or immunology assessment means travel and delay for the family, the referral letter is what keeps the question alive until they are seen.

GOOD TO REMEMBER

Ask the systemic question at the moment a psychiatric picture stops behaving like one, rather than at the end when everything else has failed. For neurosarcoidosis the trigger is a presentation crossing cognitive, psychotic and catatonic layers together. For antiphospholipid syndrome it comes from the record rather than the mental state, because the reported syndromes are the ordinary ones. For primary central nervous system vasculitis it is the neurological unease that would prompt any inflammatory cerebrovascular referral, and your written answer stops short of naming a psychiatric syndrome for it. Describe fully, attribute cautiously, route to the specialty that can settle it.

Three conditions, three different problems: neurosarcoidosis is a recognition problem, antiphospholipid syndrome an attribution problem, primary central nervous system vasculitis an evidence problem, so describe fully, attribute cautiously, and route each to the specialty that can settle it.

17 · Coeliac disease, inflammatory bowel disease and the gut-brain axis in psychiatry

Two bowel diseases reach the psychiatry examination for opposite reasons: coeliac disease for what it is associated with, and inflammatory bowel disease for how it behaves over time. Both are usually taught under the single heading of the gut-brain axis, and that heading does them a disservice, because it places an epidemiological observation and a mechanistic theory in the same sentence and lends them the same confidence. This section separates them: what the comorbidity evidence in coeliac disease states and what its design will bear, what the bidirectionality evidence in inflammatory bowel disease states and what it will not bear, what each finding changes in front of a patient, and where the mechanistic account stops being evidence and becomes a story that is merely satisfying. The examinable skill here is discrimination, not enthusiasm.

Several neighbouring topics are named here and left to the sections that own them. Neuropsychiatric systemic lupus erythematosus, neurosarcoidosis, central nervous system vasculitis and antiphospholipid syndrome each have their own home in this chapter, and their manifestations are not repeated here. The general immune-brain vocabulary of inflammatory signalling and sickness behaviour belongs to the earlier notes that own it and is used here only as a name. Steroid-induced psychiatric disturbance is owned by the routing section of this chapter and appears below only as a confound to be excluded before anything else is concluded. Delirium assessment belongs to the notes on stroke, traumatic brain injury and focal brain syndromes. Organ-adjusted prescribing, including any adjustment for malabsorption, belongs to the psychopharmacology notes, so no dose, no interaction ranking and no drug choice is stated in this section.

17.1 Three different claims wearing one name

The phrase **gut-brain axis** is used to make at least three claims that differ sharply in how well they are supported, and most of the confusion in this topic comes from sliding between them

inside a single answer.

- A **co-occurrence claim**: people carrying the bowel diagnosis carry more psychiatric morbidity than some comparison group. This is a question about counts in two populations, and it is the easiest of the three to test.
- A **covariation claim**: within a single person, the state of the bowel and the state of the mind move together over time, and each may carry information about the other's future. This needs longitudinal data and is harder to establish.
- A **mechanistic claim**: a specified biological pathway carries the influence in one direction or both. This is the claim the informal literature makes most often and the one it supports least well.

The two diseases in this section illustrate the first two claims respectively. Neither of the findings taught below is a mechanistic claim, and keeping that straight is most of the work an examiner is testing.

17.2 Coeliac disease: the comorbidity signal, and the shape of it

The finding to carry into the examination comes from a systematic review that compared people with the diagnosis against healthy controls. It **identified a significant increased risk for autism spectrum disorder, attention deficit hyperactivity disorder, depression, anxiety and eating disorders in patients with coeliac disease compared with healthy controls**. That sentence rewards slow reading, because each of its parts constrains what may afterwards be done with it.

The comparator is **healthy controls**, and that one phrase governs the whole interpretation. A comparison against healthy people cannot separate what belongs to coeliac disease specifically from what accompanies having any chronic diagnosis at all: a restrictive treatment, a lifetime of appointments, a label, and the ordinary demoralisation of managing a condition that never resolves. A candidate who says the finding demonstrates a disease-specific psychiatric vulnerability has reported a comparison that was never made. The defensible formulation is that psychiatric morbidity is raised relative to people without the disease, and that the design does not tell us whether it is raised relative to people carrying a different chronic illness of similar burden.

The outcome list is broad rather than narrow. It **spans neurodevelopmental, affective, anxiety and eating-disorder categories** rather than picking out a single syndrome, and breadth of that kind fits a general elevation of psychiatric morbidity better than it fits a specific lesion producing a specific disorder. That does not make the finding uninteresting. It makes it a reason to take a full psychiatric history in this population rather than a reason to go looking for one particular condition.

The design is a synthesis of observational comparisons. It establishes that the categories co-occur more often than expected, and it establishes nothing at all about which of them came first. The

rival explanations that survive a design of that kind are set out below, and reciting them is the cheapest mark available in this topic.

TRAP **Converting an association list into a causal list.** The review reports raised risk across several diagnostic categories against healthy controls. Candidates routinely restate this as the bowel disease causing the psychiatric disorders, then reason forward to a treatment implication that no observational synthesis can support. The finding tells you whom to assess carefully. It does not tell you what to attribute, and it does not tell you what to treat first.

17.3 What the coeliac finding actually changes in the room

Two situations arise, and they call for different work. In the first, the bowel diagnosis is already made and the person is referred with a psychiatric presentation. Here the finding does one useful thing: it removes any temptation to treat the psychiatric picture as an obvious reaction to illness that needs no further characterisation. Because the associated categories **include neurodevelopmental as well as affective and anxiety conditions**, a history that stops at mood in the last few months is too short. The developmental history, the school history and the pattern of attention and social functioning across the lifespan belong in the assessment, and they are the parts most often skipped when a medical diagnosis is already on the front of the file.

The entry in that list which most changes a consultation is **eating disorders**, because the management of coeliac disease is itself dietary exclusion. A medically required restriction and a symptomatic restriction occupy the same behavioural territory, and each can conceal the other. A person whose intake has narrowed far beyond what the bowel disease requires may be presented by the family, and even by the treating team, as unusually adherent. The clinically useful question is not whether the person avoids particular foods, which the treatment already dictates, but what the avoidance is for, what happens when it is breached, whether weight and shape have entered the reasoning, and whether the restriction has spread past the medical boundary into foods it never covered. That distinction is doable in the consultation and is where the psychiatric contribution is greatest.

In the second situation the bowel diagnosis has not been made and the temptation is to reason towards it from the psychiatric presentation. Here the section declares a limit rather than manufacturing guidance. No evidence row available to it grounds a screening rule, a case-finding threshold, an at-risk group worth testing or any performance figure for a diagnostic test, so none is printed. What survives without a row is procedural, and it is enough for most examinations: gastrointestinal symptoms, unexplained weight loss and unexplained nutritional deficiency belong to the referring physician, and the right move is a clear referral question rather than a test ordered from the psychiatry clinic.

17.4 Inflammatory bowel disease: two arrows, and the modal verb that carries them

For inflammatory bowel disease the claim to hold is different in kind. A systematic review and meta-analysis concluded that **bidirectional effects of the brain-gut axis are present in inflammatory bowel disease and may influence both the natural history of the disease and psychological health**. This is a covariation claim, not a co-occurrence claim, and it is the more

clinically interesting of the two because it concerns what happens next rather than what is present now.

Unpacked, **bidirectionality** is two arrows, and the conclusion adds a qualifier. The first arrow runs from the bowel to the mind: the state of the disease bears on psychological health. The second runs from the mind to the bowel: psychological state bears on the natural history of the disease. A complete answer names both, because naming only the first reduces psychiatric symptoms to an understandable reaction and leaves the psychiatrist with nothing to offer except sympathy, while naming only the second overstates what any single service can influence.

The qualifier matters as much as the arrows. The conclusion is that these effects **may influence** disease course and psychological health. It is a statement that bidirectional effects are present and that influence is possible. It is not a magnitude, it is not a treatment claim, and it is emphatically not a finding that treating a mood disorder alters the course of the bowel disease. That last step is the one candidates take without noticing, usually in the final sentence of an otherwise careful answer, and it converts a defensible position into an unsupported one.

TRAP **Promoting a bidirectional association into a therapeutic promise.** "The brain-gut axis is bidirectional, therefore treating depression will improve the bowel disease" reads as a natural inference and is a different claim from the one the evidence makes. Treat the psychiatric disorder because it is a disorder and it causes suffering. Do not sell it to the patient or to the gastroenterology team as disease-modifying therapy.

17.5 Working both arrows at the bedside

Bidirectionality becomes concrete as soon as it is turned into a sequence of consultation moves. Because **the effects run in both directions**, the first task is to work out where in the cycle the person currently sits. The psychiatric history on its own will leave that question open.

Begin with disease activity. Before attributing fatigue, poor concentration, disturbed sleep and reduced appetite to a mood disorder, establish with the treating team whether the bowel disease is currently active, when it was last assessed, and what the plan is. The activity measures used by the gastroenterology service are theirs to interpret and are not named or scored here; what the psychiatrist needs is the team's answer, in writing, to the question of whether this is active disease. A psychiatric formulation written without that answer is a guess about which arrow is operating.

Next exclude the treatment as the explanation. Corticosteroid-induced psychiatric disturbance is taught in full by the section of this chapter that owns it, and is named here only so that it is excluded before anything is attributed to the axis. Establish what the person is taking, what was started or stopped in the weeks before the change in mental state, and whether the timing fits. A mental state change that tracks a treatment change is a treatment problem until proven otherwise, and no amount of gut-brain reasoning displaces that.

Then characterise the illness experience rather than assuming it. Urgency, incontinence, nocturnal symptoms and the resulting narrowing of where a person is willing to go are facts that patients rarely volunteer, because they are humiliating. Asking about them directly is the part of

the interview that most often turns a generic depressive episode into a specific and modifiable set of avoidances.

RULE **Establish disease activity and treatment timing before you formulate.** In a person with inflammatory bowel disease, low mood, fatigue and poor concentration have at least four candidate origins: active disease, the treatment, the restriction the illness imposes on daily life, and an independent psychiatric disorder. These are not mutually exclusive and they demand different responses. Writing to the gastroenterology team with a specific question, and waiting for the answer, is faster than treating the wrong one.

17.6 Mechanism, and where a careful answer stops

This is the part of the topic where overreach is routine, so the boundary is worth stating plainly. The mechanistic literature on gut-brain communication is large and is written with more confidence than its evidence carries. Several families of explanation are commonly offered: immune and inflammatory signalling, the gut microbial community and its products, nutritional and absorptive consequences of the bowel disease, the effects of treatment, shared genetic or developmental liability acting on both systems, and the ordinary psychological load of a chronic, unpredictable and stigmatising illness. These are named here as categories only.

No evidence row available to this section grounds any specific mechanism in either disease, and none is asserted here. That is a declared gap, not an oversight: the two findings this section owns are an association in coeliac disease and the presence of bidirectional effects in inflammatory bowel disease, and neither identifies a pathway. A candidate asked to explain the axis mechanistically should name what is proposed, mark it as proposed rather than established, and then give what the evidence does support, which is the two claims above.

Four alternatives should be recited whenever an association in this field is offered as though it were a mechanism. The association may run in the reverse direction from the one assumed. A third factor may produce both. Detection may be differential, so that people already in one care system are more thoroughly investigated in the other. And the association may be entirely real yet wholly mediated by the burden of chronic illness, which is a psychological explanation and needs no biology at all. An answer that lists these four has demonstrated the exact competence being examined.

TRAP **Reaching for the microbiome because it is available.** A fluent mechanistic paragraph is easy to produce in this topic and is the commonest way of converting a good answer into a suspect one. Fluency is not evidence. Say what is proposed, mark it as proposed, and spend the remaining time on the two findings you can actually defend and on what they change in the consultation.

17.7 Where these two sit among the systemic presentations

It is worth placing coeliac disease and inflammatory bowel disease alongside the other systemic conditions that reach psychiatry, because they belong to a different column. The autoimmune conditions taught elsewhere in this chapter present as direct neuropsychiatric manifestations, in which the illness process itself produces the psychiatric picture and the diagnostic task is to recognise it. The two bowel diseases here are not taught that way. Their psychiatric relevance is

an elevated burden of common psychiatric disorders in coeliac disease and a bidirectional relationship with disease course in inflammatory bowel disease, which is a question of comorbidity and prognosis rather than of a direct neuropsychiatric syndrome.

The systemic autoimmune comparator plate sets each systemic presentation in this chapter beside its discriminating feature, so the conditions can be told apart in the moment rather than merely listed. Read it as a routing aid: the discriminating feature is what should redirect the assessment when it is present. The two entries relevant here are the ones whose discriminating feature is not a neuropsychiatric sign at all, and noticing that asymmetry is the point of putting them on the same plate.

That plate is printed with the neuropsychiatric lupus section, where it places these two conditions beside the other four, and it is the fastest way to revise the discriminating features together.

17.8 What a complete answer contains

A structured answer on the gut-brain axis contains five things in this order: the definition with the three claims separated; the coeliac finding stated with its comparator, which is the part most answers drop; the inflammatory bowel disease finding stated with both arrows and the qualifier intact; the design limitations, meaning direction, shared liability and differential detection; and the consultation consequences, meaning a lifespan history in coeliac disease, the required-versus-symptomatic restriction question, and disease activity and treatment timing established before any psychiatric attribution in inflammatory bowel disease.

QUESTION	COELIAC DISEASE	INFLAMMATORY BOWEL DISEASE
Type of claim the evidence supports	Co-occurrence against a comparison group	Covariation over time, in both directions
What is stated	**Increased risk of autism spectrum disorder, attention deficit hyperactivity disorder, depression, anxiety and eating disorders against healthy controls**	**Bidirectional brain-gut effects are present and may influence disease course and psychological health**
What is not stated	Direction of causation, disease specificity against another chronic illness, any mechanism	Any magnitude of effect, and any claim that psychiatric treatment alters the natural history
First move in the consultation	Full lifespan psychiatric history, developmental included, and the required-versus-symptomatic restriction question	Establish current disease activity and treatment timing with the treating team, in writing
Commonest error	Reading the association list as a causal list	Promising disease modification from psychiatric treatment

IN INDIA

No Indian prevalence figure, service audit or cost estimate for either disease appears in this section's evidence rows, so none is printed; the shortfall is declared rather than filled. What can be said is organisational. The gastroenterology service is usually in a different building from psychiatry and often in a different institution, so the liaison here is a written one, and it should carry a specific question rather than a request for an opinion: is the disease active, what changed in the treatment and when, and what is the plan for the coming months. The dietary restriction that treats coeliac disease is negotiated inside a shared kitchen, with food prepared for the household at once, so an adherence discussion aimed only at the patient misses where the decisions are actually made. And where a person travels a long distance to a tertiary clinic, align the psychiatric and gastroenterology reviews to the same visit wherever the services can manage it, because a review that needs a second journey is a review that will not happen.

GOOD TO REMEMBER

In coeliac disease, widen the history: the associated categories reach into development and into eating behaviour, so a history that covers only the last few months of mood is the wrong shape. In inflammatory bowel disease, narrow the question: find out whether the disease is active and what the treatment has been doing before you attribute anything to the axis. In both, treat the psychiatric disorder because it causes suffering and is treatable, and resist the pull to sell it to the patient or to the referring team as a way of changing the bowel disease.

MNEMONIC

Widen, then narrow

Widen for coeliac disease, because the associated conditions cross the lifespan and the eating-behaviour boundary. Narrow for inflammatory bowel disease, because the useful question is which arrow is operating right now, and that is answered by disease activity and treatment timing rather than by mechanism.

Three, not one	Comorbidity, not cause	Possible, not measured	Proposed, not established
CLAIMS INSIDE THE ONE PHRASE	WHAT THE COELIAC REVIEW SUPPORTS	WHAT THE QUALIFIER LEAVES OPEN	HOW MECHANISM MUST BE MARKED

Coeliac disease gives you a comorbidity fact and no direction; inflammatory bowel disease gives you a direction in both senses and no magnitude. Everything beyond those two sentences that sounds mechanistic is a proposal, and saying so out loud is what separates a careful answer from a fluent one.

18 · Burn injury and its psychiatric course

Burn injury is the one medical event in this chapter where the psychiatric question runs forwards and backwards: what the injury has done to the patient, and what the patient may have been doing when it happened. Everything examinable here follows from holding those

apart. This section installs a timeline, because each phase of burn care asks a different question, then a question that sits outside the timeline because it is asked at presentation and concerns the moment before the flame. It also states what this text can source and what it cannot, since burns psychiatry is taught nowhere else in these notes.

18.1 The acute phase: the ward you are called to

The referral comes from a surgical or burns team, and psychiatry walks into an environment organised around fluid balance, wound and airway. The mental state you are asked about is produced by several processes at once, and the acute-phase skill is separating them rather than settling on a label. Distress belonging to the event, distress belonging to the pain, a disturbance of consciousness and the effects of what is given for the pain are all live in the same bed at the same hour, and each responds to something different.

What was sought here, and not obtained in verifiable form, was a quotable statement from a burns or trauma source able to carry it that psychological shock, **dissociation** and **acute stress disorder** may occur in the acute phase after burn injury. State the reason exactly, because it changes what the gap means: of the burns records assembled for this section, several hosts refused automated retrieval, one address no longer resolved, and one quotation could not be found in the document actually served. Each of those is a fact about retrieval, none about burn patients. The claim, if made, would rest on the burns and trauma literature, where acute-phase presentations after burn injury are described from real cohorts, and which psychiatry reads rather than owns.

Without the figure, you describe. Acute distress after a burn injury is observable at the bedside and needs no prevalence attached. Record the phenomenology in the patient's own words, date it against the injury and against each operation, and note whether it fluctuates. That record lets a later clinician settle the question you could not, and it beats a confident figure from nowhere.

One process must be named and handed straight over. Disturbed consciousness, disorientation and fluctuating attention occur in burn patients, and the acute burns bed is one place a psychiatrist meets them. This chapter states that they occur and stops there. Assessment, motoric subtypes and management belong to the Stroke, TBI, delirium and focal brain syndromes notes. The routing is deliberate: a second delirium pathway written inside a burns section is how sibling chapters come to disagree, and a disagreement about an acutely unwell patient is a safety defect rather than an editorial one.

RULE **In the acute phase, ask which process, not which label.** Acute distress, pain, disturbed consciousness and drug effect co-occur in a burns bed and respond to different things. The examinable act is separating them at the bedside and routing each to its owner: the disturbance of consciousness out of this chapter, analgesia to the treating team.

18.2 Pain, and the part of it that is yours

Pain is the largest psychiatric problem of the acute phase and the one most easily handed back whole. Burn pain is unusual in being renewed deliberately: dressing changes, debridement and physiotherapy each produce a further episode of severe pain at a time the patient can predict, so

anticipation becomes part of the pain. That anticipatory limb is psychological and sits inside your remit whether or not the analgesic chart does.

The sourced statement available here is modest and useful. In a systematic integrative review of psychotherapeutic interventions for burns patients, **relaxation breathing techniques are an intervention to be considered for pain and anxiety relief in the care of burns patients**. The pairing is the teaching: the intervention is offered for pain relief and for anxiety relief together, which matches what the ward presents, a patient whose pain and whose fear of it have become one complaint by the third dressing change.

Its practical virtues belong in an answer: it needs no equipment, costs nothing, can be delivered by nursing staff once taught, and can be rehearsed with the relative at the bedside. It is an adjunct, and an adjunct sits alongside analgesia rather than in place of it. Analgesic prescribing stays with the treating team and this chapter prints no drug and no dose. Psychotropic selection and any organ-related adjustment belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes, which this chapter adds no further version of.

TRAP **Handing pain back whole.** Candidates write that analgesia is the surgical team's business, which is true of the prescription and false of the problem. Anticipatory anxiety before a predictable painful procedure is a psychological mechanism with a psychological adjunct, and an answer that omits it surrenders the part of burn pain psychiatry was consulted about.

18.3 Sleep, and where its causes sit

Sleep fails early after burn injury, and it fails for reasons that come from more than one place. What was sought here, and not obtained in verifiable form for the retrieval reasons already given, was a quotable statement that poor sleep in burn patients reflects psychological factors and the hospital environment together. That claim would rest on the burns and trauma literature, and on the ward-based sleep work within it.

The gap costs little at the bedside, because the instruction underneath it survives intact. Take the sleep history so that the sources separate: on one side what the ward is doing, meaning light through the night, noise, observations, position and procedures timed for the day list; on the other what the mind is doing, meaning intrusive re-experiencing, nightmares, worry about appearance and money, and the anticipated dressing change. Ask about each explicitly, because a patient asked only whether they are sleeping badly will say yes and tell you nothing actionable. Then correct the environmental limb first, since a hypnotic prescribed into a bright and noisy bay treats a cause that was never in the patient.

18.4 The subacute phase: appearance and disfigurement

As the patient stabilises, the question changes from survival to what the body now looks like and what that will mean for the life attached to it. Keep these apart. **Disfigurement** is a visible difference in the body, assessable by anyone looking at it. **Appearance concern** is what the person makes of that difference, assessable only by asking them. An assessment that measures the burn and infers the distress has measured the wrong variable, and it will be wrong in either direction.

The sourced finding comes from a multi-centre prospective cohort of adults hospitalised after burn injuries. **Increased psychological flexibility during hospital admission after a burn injury predicts lower appearance concerns over time, although not over the effect of baseline appearance concerns.** Each part carries an examinable point. The exposure was measured during the admission, so the acute admission is a measurement window and not only a treatment window, and the outcome came later, so an acutely measured variable travels forward into a later-phase result. That is why this section is a timeline rather than a list of complications. And the qualification is what an examiner will press: the prediction does not hold over and above baseline appearance concerns, so what the person already thinks about their appearance carries the prognostic weight.

The clinical consequence is direct. Ask about appearance concern during the admission, because it carries the prognostic weight and costs one question, and record it so that later change can be read against something. Be careful, though, describing any intervention aimed at flexibility, because the study supporting the association is the study that qualifies it.

TRAP **Reading a predictor as a treatment target.** A variable that predicts an outcome, then loses its prediction once a baseline is controlled for, is telling you where to look rather than what to change. Candidates keep the first clause of that sentence and drop the second, and the dropped clause is why the finding was reported as it was.

Social re-entry also begins in the ward rather than at the front door, and the work is concrete: rehearsing the first meeting with people who have not yet seen the scars, deciding what the patient will say when a stranger asks, and preparing the family, whose answers to the same question set the tone the patient has to live with.

PHASE	THE QUESTION IT ASKS	WHERE IT IS ANSWERED
Acute	Which process is producing this mental state, and what does each need?	This section, for acute distress and for the psychological limb of pain and sleep.
Acute	Is this a disturbance of consciousness, and how is it assessed and managed?	The Stroke, TBI, delirium and focal brain syndromes notes. Named here, never taught here.
Subacute	What does this person make of the disfigurement, and how closely are they followed?	This section, with its prevalence gaps declared.
Any phase	Which psychotropic, and how is it adjusted for organ impairment?	The Psychopharmacology notes. This chapter adds no further version.
Any phase	Does this patient have capacity to refuse the dressing change?	The Mental health legislation notes. Capacity is decision-specific and time-specific; the test is taught there.
At presentation	How did this burn injury happen, and was it an act?	This section, as an antecedent question.

18.5 Depression and post-traumatic stress: what is missing, and what holds

Here the declared gaps bite hardest. What was sought first was a citable prevalence of post-traumatic stress disorder after burn injury, as a proportion of a defined cohort at stated follow-up points inside the first year. What was sought second was a pooled estimate linking burn extent, as a proportion of **total body surface area** above a stated threshold, to an increased risk of depression expressed as a risk multiple. Neither was obtained in verifiable form, so this text supplies neither the proportions, nor the threshold, nor the multiplier. Each would rest on the burns and trauma literature, on primary follow-up cohorts for the first and a quantitative synthesis of such cohorts for the second.

Now the part that carries marks. A missing figure is not a negative finding, and converting one into the other is a larger error than the one it replaces. To say that post-traumatic stress after burn injury is uncommon, or that burn extent is unrelated to later depression, asserts far more than a failed retrieval can support. The correct sentence is short: the figure exists in the burns literature, I cannot state it from a source, and it does not change what I do with the patient in front of me. It does not, because prevalence tells you how hard to look, never whether this patient has the syndrome. That is settled by assessment, applied during admission, at the transition out of acute care, and across the long rehabilitation tail when the operations have stopped and the attention has thinned.

What holds, sourced and directive, is what to do with a patient who screens badly. European practice guidance for burn care states that **patients with an abnormal psychological profile, including suicidal ideation, should be given special attention, adequately monitored and regularly followed by appropriate mental health professionals**. Notice its shape. It is written around a profile rather than a diagnosis, so it is actionable on the day of the assessment and before any diagnostic conclusion is reached. It names obligations of distinct kinds, attention and monitoring and follow-up, which are commonly collapsed into one vague plan. And it names suicidal ideation inside the profile, which is the hinge into the last part of this section.

TRAP **Turning a missing number into a negative finding.** The candidate who cannot recall a prevalence says the association is weak or the disorder is rare. That is a categorical claim built on a failure of memory, unsafe in exactly the domain where being wrong costs a life. Name the gap, name whose literature holds the figure, and state what you do anyway.

18.6 Rehabilitation, and the return to work

Rehabilitation after a major burn injury is measured in years. Scars mature slowly, contractures develop and are released, and each new operation re-opens the acute and subacute questions in a patient who had begun to leave them behind. The consequence is unglamorous and examinable: psychiatric follow-up must outlast surgical follow-up, and the plan should say so in writing, because attention thins once the wounds are closed and it is then that appearance concern and social re-entry become the whole of the patient's day.

A further gap is declared here. What was sought was a comparison of work-reintegration outcomes between burn survivors whose injury happened at work and survivors injured

elsewhere, with the direction of the difference stated. It was not obtained in verifiable form, and it would rest on the burns rehabilitation and occupational health literature, where such cohorts are followed and outcomes defined well enough to compare.

The instruction underneath it survives the gap: ask where the burn injury happened, because the answer reshapes the rehabilitation problem. When the injury occurred at work, the site of the injury and the site of return are the same place, and that single fact brings in fear of the machine or the flame, the employer's response, blame, compensation and colleagues who witnessed it. A domestic injury raises its own version, since the kitchen is also a place the patient must return to. In Indian practice much work is informal, so returning to work may mean returning to a stove, a field or a roadside stall, with no graded return to arrange and financial pressure that overrides advice framed as a preference.

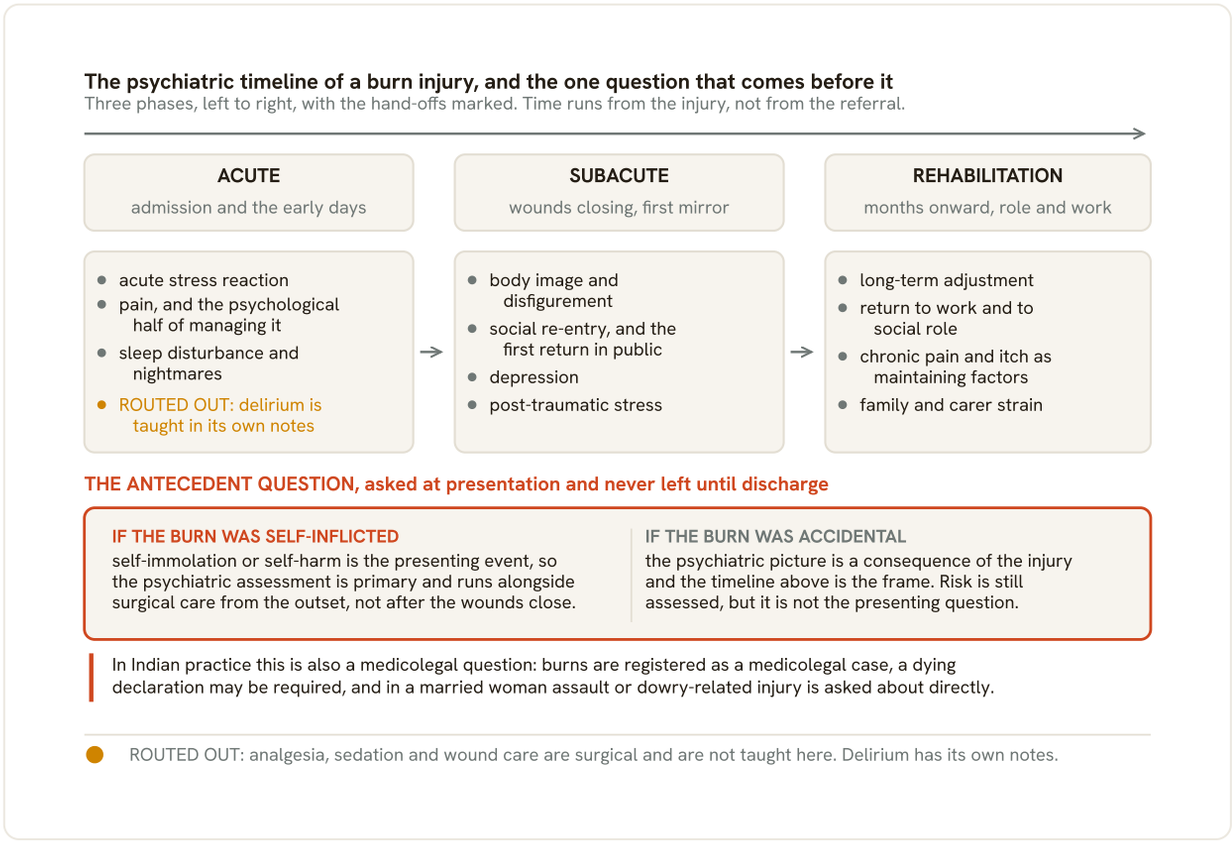


Fig 40.10 The psychiatric course of a burn injury across three phases, with the antecedent question set apart from the timeline because it is not part of it. This is the only plate in the series that teaches burns, so it is built to carry the whole construct rather than illustrate one corner of it. The phases are not merely descriptive: each has a characteristic failure, which is treating acute distress as a disorder before the wounds have closed, missing the subacute period in which body image and the first return in public do most of the psychological damage, and discharging at wound healing as though the rehabilitation phase were somebody else's problem. The block below the timeline is the discrimination that changes management, because a self-inflicted burn makes the psychiatric assessment the primary one from the outset rather than a consequence to be picked up later, and in Indian practice the same question carries a medicolegal weight that an accidental injury does not. From memory you should be able to name the phase-specific problems in order, say where the delirium pathway hands off, and state why self-harm is asked about as an antecedent at presentation rather than as an outcome at discharge.

18.7 The antecedent question, asked at presentation

Everything above treats psychiatric problems as consequences of burn injury. There is a question that runs the other way, asked at presentation rather than at discharge: was the injury itself an

act, and was psychiatric disorder present before the flame rather than after it. That discrimination changes management, and it is the one most often lost when a burns referral is treated as routine liaison work.

It changes management because consequence-work and antecedent-work run on different clocks. Consequence-work is rehabilitative and can be paced. Antecedent-work is risk-work and cannot be: if the burn injury was self-inflicted, the patient is already inside a high-risk period, the method remains available on discharge, the collateral history has to be taken while relatives are still in the hospital, and the discharge destination becomes a psychiatric question. A service that asks how the burn happened only when the wounds are healing has asked its most important question several weeks late.

The Indian relevance is direct, because self-immolation presents to burns units here rather than to psychiatric services, and because the Indian psychiatric literature has examined it. **An Indian report investigated four cases of attempted self immolation, with the psychiatric evaluation done according to DSM-III-R criteria.** Be exact about what that establishes. It is a small clinical series with formal diagnostic assessment applied to it, so it establishes the construct as one Indian psychiatry has studied with criteria rather than impression. Its denominator is small, and using it as though it described a national picture would misuse it.

The same literature shows that self-immolation is heterogeneous rather than a single clinical entity. In an Indian psychosocial study, **22 young people, mostly students, had indulged in this act to protest against the decision of the Government of India**. A public and collective act of protest is a different phenomenon from a private act of self-harm during a domestic crisis, and a formulation built for one will misdescribe the other. So the clinical question is simple to ask and easy to skip: what was the act for. The same physical burn injury can follow any of these.

- A suicidal act, after which the patient is already inside a high-risk period and the method remains available on discharge.
- A public and collective act of protest, which is the phenomenon the Indian psychosocial study above describes.
- An impulsive act inside a family quarrel, a private act during a domestic crisis rather than a public one.
- An accident, where the injury was not an act, though the question of psychiatric disorder before the flame still stands.
- An assault by another person, which this section names as an antecedent and does not follow further.

Those demand different assessments and different collateral sources, which is why the question belongs at presentation rather than in the discharge summary.

Where the assessment finds a self-inflicted injury, ongoing suicidal ideation, or a disturbed psychological picture of any kind, the sourced instruction above applies directly and is worth quoting rather than softening: **special attention, adequate monitoring and regular follow-up**

by appropriate mental health professionals. In a burns unit those obligations have concrete forms. Attention means your assessment happens at presentation, in the surgical bed, written where surgical staff will read it. Monitoring means the nursing observations reflect a psychiatric risk as well as a wound. Follow-up means a named appointment attached to a named person, since a patient who leaves with a dressing clinic date and a vague psychiatric intention will attend the first and lose the second. Where such a patient then refuses dressing changes, capacity becomes the question, and capacity is decision-specific and time-specific; the statutory test is taught in the Mental health legislation and forensic psychiatry notes.

IN INDIA

The burns unit is usually where this psychiatry happens, because the patient occupies a surgical bed and the psychiatric team travels to them. Your assessment competes with a dressing round, your notes are read by surgeons, and the family outside is your collateral source and the carer who will do the dressings at home. Distance and cost shape everything after discharge, so a plan written for a family travelling a long way looks different from one written for a city outpatient. Note what this text could and could not verify. The acute-phase, prevalence and work-reintegration claims were not obtained in verifiable form here, which is a limit of this text rather than a statement about the Indian record. What was verified is Indian, the self-immolation work, and that tells you where Indian psychiatry has looked hardest at burns.

GOOD TO REMEMBER

Ask how the burn injury happened before you ask how the patient is coping with it, because the first question sets the urgency of everything else. In the acute bed, separate the processes and route the disturbance of consciousness to the Stroke, TBI, delirium and focal brain syndromes notes rather than managing it here. Treat anticipatory pain as your work and offer the psychological adjunct alongside the analgesia the team prescribes. Ask about appearance concern during the admission, because it is the early variable that carries the later weight. Where a figure is missing, name the gap and the literature that owns it, then say what you will do anyway.

Antecedent first	Which process, not which label	Asked in the admission	A named gap, not a negative
ASKED AT PRESENTATION, NOT AT DISCHARGE	WHAT THE ACUTE BURNS BED ASKS OF YOU	WHEN APPEARANCE CONCERN CARRIES THE LATER WEIGHT	WHAT A PREVALENCE YOU CANNOT SOURCE BECOMES

Burns psychiatry runs in two directions: ask at presentation whether the injury was an act, then follow the timeline of what the injury has done, separating process from label in the acute bed, treating anticipatory pain as your work, asking about appearance concern during the admission, and naming a missing figure rather than converting it into a negative finding.

Consolidate and apply

19 · Worked vignettes

Every referral in this chapter asks one question in a different accent: is this the medical illness, is this the treatment for it, or is this a psychiatric disorder that would have arrived anyway. The sections before this one supply the content for that triage. What they cannot supply is the experience of running it against a referral somebody else has already answered, where all three possibilities are live at once and the answer is sometimes that the question belongs to another chapter.

Each vignette opens with the referral as it would arrive, works the reasoning in the open, and closes with what an examiner is testing. Where a vignette reaches a specific taught elsewhere, it names the owning section and stops. Two end without a treatment, because the answer is a hand-over. One ends without a fact, because the fact was sought and not established.

19.1 A decompensated heart and a request for an antidepressant

The referral. From a cardiology ward. A man in his sixties, readmitted within months for decompensated heart failure. Low mood, poor appetite, sleeping badly, no motivation, will not participate with physiotherapy. Please see and start an antidepressant.

The reasoning. The failing circulation can generate every item on that list. Fatigue, appetite change, broken sleep, slowed movement and withdrawal from effort follow equally from the haemodynamic state, from the treatment burden, or from a mood disorder, so the count will be high whichever is true and says nothing about which of the three is operating.

What separates them is **order in time**. Put the cardiac events, the regimen changes and the mood changes on one timeline. If low mood arrived with the decompensation and lifted as it settled, that is a person responding to a frightening event. If the breathlessness and the swelling improved while interest, outlook and self-reproach stayed where they were, the mood state has separated from the haemodynamic state and behaves as an illness in its own right. The second lever is pleasure in things costing no physical effort, because a man too breathless to reach the gate can still say whether a visit gives him anything.

Keep the treatment limb open: the regimen is itself a psychological load, and admission interrupts whatever he took at home. Then look again at what was asked. The note has gone straight to a prescription, and agent choice in established cardiac disease is not this chapter's to give: it belongs to the psychopharmacology notes, and prognosis with safe antidepressant choice after myocardial infarction to the merged organ-interface routing section.

TRAP **Accepting the triage the referral has already performed.** The note has decided that this is a psychiatric disorder and that the intervention is a tablet. A consultation answering only the question asked leaves the attribution unmade, and reads later as agreement.

What the examiner is testing. Whether you attribute symptoms rather than count them, and whether you can be useful to a cardiology team without writing a prescription.

19.2 Breathless at night, and already told it is anxiety

The referral. A woman in her forties with obstructive airways disease is referred to the psychiatric outpatient clinic with episodes of breathlessness, palpitations and terror that wake her from sleep. The letter states that her airflow is stable in clinic, that the attacks are anxiety, and asks for treatment of panic.

The reasoning. The whole difficulty sits in the letter's middle clause. Breathlessness and panic recruit the same body and the loop closes both ways: airflow limitation produces fear, fear produces hyperventilation, hyperventilation produces breathlessness. The features shifting the probability, the temporal shape of the episode, its relation to exertion or rest, and the content of the thought at the peak, are taught by this chapter's section on obstructive airways disease, asthma and chronic respiratory failure. This vignette is about what a clinician does when they shift the probability without settling it, which is most of the time.

Two things follow. The commonest correct answer here is both rather than either, and each condition worsens the other, so treating one and closing the case leaves the disability where it was. And the errors are not symmetrical: calling a panic attack an exacerbation costs an unnecessary course of treatment, while calling an exacerbation a panic attack can cost a life. Note too what the letter's premise establishes. Stable airflow in a daytime clinic is a statement about the clinic, not about three in the morning, and nocturnal episodes have nocturnal explanations a daytime measurement cannot exclude.

One contribution is concrete, not diagnostic. Coexistent anxiety and depressive symptoms are routinely used to filter patients out of the rehabilitation that treats the deconditioning, and the guideline position blocking that is stated, in its exact wording, by the owning section.

RULE **When the triage cannot be completed in the room, complete the safe half and record why.** Treat the airway first, interrogate the psychiatry once the episode has settled, and write down which half you answered today. A note committing to a label it cannot defend is more dangerous than one stating what it could not decide.

What the examiner is testing. Whether you accept a triage handed to you in a letter, which direction of error is dangerous, and whether you can offer both with an order of work.

19.3 Confusion after transplantation, and what the psychiatrist does not own

The referral. A call from a transplant unit, late evening. A man in his thirties, some weeks after receiving a kidney from a living relative, has become confused, does not recognise the ward and has had a witnessed seizure. He is on maintenance immunosuppression and has had pulsed treatment for suspected rejection in recent days.

The reasoning. Several sufficient causes are present at once, so the examinable skill is **sequencing rather than selection**. He may be septic while barely febrile, metabolically deranged from graft dysfunction, hyponatraemic or hypoxic, on more than one agent with recognised

neuropsychiatric toxicity, and carrying the load of a graft he believes he could lose. An opinion that names the drug and stops has chosen the interesting cause rather than the first one.

The sequence that survives scrutiny puts reversible and dangerous causes first, keeps the immunosuppressants permanently on the list rather than at the end of it, and treats the psychological formulation as concurrent rather than as a diagnosis of exclusion. One technique is worth taking away whole: the drug history here is a sequence, not a list. What was given, when relative to the event, and what was given alongside it carry the information, because the labelled neurotoxicity signals in this class attach to changes of regimen and to combinations. A chart transcribed as current medicines has discarded the structure that made them reportable.

The boundary, which is the point of this case. Three of the four things the team wants are not yours here. Detection, rating and management of the acute confusional state, including any drug used for it, are taught once, with the stroke, traumatic brain injury, delirium and focal brain syndromes notes. Which co-prescribed agent alters exposure to which immunosuppressant belongs to the psychopharmacology notes. The named neurotoxic syndromes of the individual agents belong to this chapter's transplant and immunosuppressant section. You hand over a dated mental state, a timeline against the drug chart, a statement that the immunosuppressant stays on the differential, and the pathway for each remaining question.

RULE **A referral you cannot answer completely is answered by naming who answers each part.** Routing is a clinical act, not an evasion. It moves a question to somebody who can settle it, and it beats a partial version from memory that commits the receiving team to something you never verified.

What the examiner is testing. Whether you sequence rather than select, and whether you can be useful while declining most of what was asked.

19.4 Two questions from a surgical ward, and three different owners

The referral. The surgical team telephones the day after major abdominal surgery on a man in his seventies. He is agitated at night, pulling at his lines, settled by day. They ask two things in one breath: does he have capacity to refuse the drain, and is this postoperative cognitive dysfunction.

The reasoning. Separate the questions before answering either, because one undifferentiated account of postoperative confusion loses marks in both directions at once. Capacity first. One property does most of its work here: capacity is decision-specific and time-specific. The question is never whether this man has capacity as a standing attribute, but whether he can make this decision at this hour, and the perioperative window supplies an unusual density of moments in which the answer moves: pain, disturbed sleep, premedication and the hours after anaesthesia all sit inside it. An opinion formed on the morning round does not transfer to a conversation that evening. Everything else is elsewhere: the statutory test and the procedure following a finding of incapacity are taught with the mental health legislation notes and are neither restated nor cited by provision. You describe functional ability for the named decision at the named time; the decision belongs to the treating team.

The second question turns on an observation nobody on the ward tonight can make. A cognitive change is a relation between two observations, so the first thing to establish is not what is wrong with him tonight but whether anybody measured him before the operation. If not, the trajectory question cannot be settled on this ward at all, and the honest note says so, manages the acute presentation through the chapter owning it, and arranges the review after discharge.

TRAP **Recording a global capacity verdict.** A note stating that the patient lacks capacity, with no decision named and no time attached, is unusable by the team receiving it and indefensible if later examined. Worse, it travels: written once, it is quoted for months across decisions it was never about.

What the examiner is testing. Whether you separate two questions that arrived as one, and whether your first move is to ask about the measurement taken beforehand.

19.5 A cognitive complaint after treatment ended, and a question the evidence does not answer

The referral. A woman in her fifties, referred by her family doctor some months after finishing chemotherapy. Back at work on reduced hours, she cannot hold several things in mind, loses words she knows well, and works far more slowly than before. She asks whether this will get better, and when.

The reasoning. Take the complaint seriously and decline its title. The word patients use for this names a cause and, in naming it, stops the clinician looking. Run the triage instead.

- The illness limb: fatigue, disturbed sleep, pain and the endocrine consequences of treatment all fail exactly these effortful functions.
- The treatment limb: the cytotoxic agents are one contributor inside a bundle that also holds the surgery and the anaesthesia.
- The independent-disorder limb: an untreated depressive episode degrades attention, retrieval and speed, and is the most treatable thing on the list.
- A fourth possibility peculiar to this population: something still evolving underneath, so whether the picture is static or moving is asked deliberately.

The declared gap, which is her actual question. She has asked for a trajectory. What can be said is that a change of this kind has been reported as still detectable at the end of the interval examined, and that an interval describes how far the looking reached, not where the problem ended. A statement of the course beyond that point would rest on longer follow-up in the oncology and neuropsychology literature, and none was obtained for these notes in a quotable form. The absence describes what these notes opened, is not evidence that the course is unknown to the field, and must not be converted into reassurance or into catastrophe.

What you do without it is concrete. Treat the mood, sleep and pain degrading her cognition today. Characterise the deficit well enough for her workplace to be adjusted around it, whether or not it reverses. And set a review point, because the moment the oncology course closes is the moment nobody is reviewing cognition.

RULE **A declared gap is a complete answer and it has three parts.** State what was sought, on whose authority the missing claim would rest, and what you do without it. An answer carrying only the first two is a shrug; one supplying a remembered figure is worse, because the person receiving it cannot correct it.

What the examiner is testing. Whether you convert an observation window into a prognosis, and whether you can still leave the room with a plan after saying that something is not established.

19.6 Psychosis in an established autoimmune disease, where the differential is the whole answer

The referral. A woman in her twenties with a long-standing systemic autoimmune disease is admitted under medicine with a few days of persecutory beliefs, hallucinations and badly disturbed sleep. Her disease has recently been treated more intensively. The team has sent a serology panel and asks whether this is the disease affecting her brain.

The reasoning. The referral asks for an attribution and offers a laboratory result as the route to it, and neither is available in the form requested. Serology establishes that she has the disease. It does not establish that the syndrome in front of you was produced by it, and a patient with unambiguous systemic disease can equally have an infection, a drug effect or an unrelated psychiatric illness. Attribution here is a reasoned judgement rather than a measurement, and it is the substance of the case rather than its preliminary.

So the answer is an ordered list of what must be displaced before the disease is credited. Infection first, and staying first, in a person immunosuppressed by both the illness and its treatment. Metabolic derangement and blood pressure next, particularly where there is renal involvement. Then the drug chart, read with the dates set against the onset, because the treatment of the disease and the cause of the psychiatric picture can be the same agent; that construct is owned by the routing section of this chapter and by the endocrine and metabolic notes, and is named here only as a direction of travel. Then a primary psychiatric illness with a history predating the systemic diagnosis, and the ordinary response of a young person to a relapsing, visible illness with an unpleasant treatment.

The order matters because of where the answers lead: one path escalates immunosuppression and the other reduces it, so an attribution made the wrong way round is a harm rather than an inelegance. The manifestations and the treatment recommendation with its restrictions are taught in full by this chapter's section on the condition. Give the attribution as a judgement, with the observations it rests on and what would change it. A judgement can be revised without anybody losing face; a finding is very hard to withdraw.

TRAP **Treating a positive result as the answer to the referral question.** The panel speaks to the disease, never to the syndrome, and answering with it skips the contested part of the case. It is also how a commoner and more reversible cause is missed while treatment is escalated in the wrong direction.

What the examiner is testing. Whether you recognise that the difficulty is attribution and not treatment, and whether you write a judgement rather than a verdict.

19.7 A burns unit referral, where the first question is how the injury happened

The referral. A burns unit asks for review of a man in his twenties, several weeks into an admission after a flame injury. He is described as low, withdrawn, sleeping badly and increasingly reluctant to cooperate with dressing changes. The unit would like something started.

The reasoning. The triage runs as it does everywhere in the chapter, but the order is unusual, because one question sits underneath all three limbs and gets harder to ask with every day it is left. Ask first how the injury happened. A burn may be an accident, an assault or a deliberate self-inflicted act, and the risk formulation, the conversation with the family, the safeguarding position and the discharge plan differ completely across those possibilities. Ask early, ask directly, ask the patient and whoever brought him, and record the accounts separately, because they harden with retelling and a unit that has settled on an accident will not re-open it.

Then work the ordinary causes, which this setting stacks. Pain, and specifically the anticipatory dread of a procedure repeated on a schedule, which is a distinct and treatable phenomenon rather than a general lowness. Sleep, degraded by the injury, by the analgesia and by a unit lit and busy through the night. Acute confusional states, which occur here and route out entirely to the notes owning delirium, including their bedside assessment and every drug used for them. And an independent depressive episode, which this setting makes easy to explain away, since comprehensibility is not a diagnostic exclusion.

The word uncooperative is a **referral word**, and it repays unpicking, because it names a behaviour and hides its reason. Reluctance to face dressings after a course of them is a learned response to a procedure that has reliably hurt, and it answers to analgesia, prediction, a share of control and specific procedural anxiety work, none of which is persuasion and none a tablet. Beyond the admission sit appearance, the reactions of others and return to work, all of which belong to this chapter's burns section.

TRAP **Reading a refusal of dressings as either a character problem or a depressive symptom.** Both readings skip the mechanism and both produce the wrong intervention, one punitive and one pharmacological. The behaviour is usually specific to the procedure, which is exactly what makes it treatable.

What the examiner is testing. Whether you ask how the injury happened before anything else, and whether you can decompose an unhelpful referral word into treatable parts.

The cases share a shape: the referral names a symptom and requests a treatment, while the real question is about order in time, the rest of the chart, or ownership.

THE REFERRAL AS WRITTEN	THE QUESTION UNDERNEATH IT	WHO OWNS THE ANSWER
Low mood in heart failure, start an antidepressant	Did mood move with the circulation, or separately	Attribution here; agent choice, psychopharmacology
Her attacks are anxiety, treat panic	What drives tonight's episode, and which error is dangerous	The obstructive airways and respiratory failure section
Confused and fitting after transplantation	What was given, when, and alongside what	Sequence here; the state itself, delirium; interactions, psychopharmacology
Capacity to refuse, and is this cognitive dysfunction	Which decision, at what hour, and was he measured beforehand	Decision-specific and time-specific here; the test, legislation
Will my concentration come back	What degrades it today, and over how long a window was it observed	Function here; the findings, the chemotherapy and interferon section
Is this the disease affecting her brain	What must be displaced before the disease is credited	Attribution here; manifestations, the owning section
He is not cooperating, start something	How did the injury happen, and what is he avoiding	Circumstances here; the sourced material, the burns section

Four moves sit under these cases, and not one of them is a fact to recall.

Re-opened, not accepted
THE TRIAGE THE REFERRAL ARRIVES WITH

Attributed, not counted
WHAT MAKES A SYMPTOM LIST MEAN ANYTHING

Named, not attempted
THE PART OF THE REFERRAL THAT IS NOT YOURS

Sought, not remembered
A FACT THESE NOTES COULD NOT ESTABLISH

The first move is the one the other three depend on, because a triage accepted at the door decides the answer before any of them is run.

IN INDIA

Several features of Indian practice change how these cases run, and none is a figure, because this section holds none. Referrals arrive late and pre-decided, so the first act is usually to re-open a triage somebody else closed. A relative is almost always present, which makes collateral history easy to obtain and the funding question answerable at the bedside. Follow-up happens by travel rather than by schedule, so a plan whose active ingredient is a review appointment fails quietly. And how a burn happened carries legal consequences from the day of admission.

GOOD TO REMEMBER

Before answering any referral in this chapter, ask three questions in order. What has changed, in what order, against what baseline. What else on the chart or in the admission could produce the same picture. And which part of this is not mine to answer here. The first settles the attribution, the second protects the patient, the third makes the note usable.

Run the triage rather than accepting the one in the referral letter, attribute by order in time rather than by symptom count, name the owner of every question that is not yours, and where a fact was not established, say what was sought, whose claim it would be, and what you do anyway.

20 · Consolidation and quick review

In the last week before the examination the useful question is not what this chapter contained but what it installed, and it installed one triage, one habit of naming owners, and one way of saying what is not established. Every section is a worked instance of those three. This consolidation restates them in the smallest form that survives questioning: the triage and the evidence deciding each limb, a grid of the features separating the conditions the chapter teaches, the boundaries stated as a competence rather than an apology, the shapes an answer should take, and what to read if there is an hour left. Read it as a set of moves rather than a summary, because content cannot be compressed without becoming the thin second copy the chapter argued against, whereas reasoning can.

20.1 The three-way question, and what each answer is made of

The chapter opens with a triage and never leaves it. A psychiatric presentation in a medically unwell patient is being produced by the illness acting on the brain, by the treatment for that illness acting on the brain, or by a psychiatric disorder that is the patient's own and has arrived in a body that happens to be unwell. Those are the three limbs, and the chapter is training in choosing between them and in saying why.

Two properties matter more than the categories. The limbs are not exclusive, and more than one is usually running at once, so the working question is never which limb is true but which is doing the most damage now and which, if missed, is dangerous. Each is settled by a different kind of evidence, which is why an answer naming a syndrome without placing it against a limb has skipped the examined step.

LIMB	WHAT SETTLES IT	WHAT IT OBLIGES	CHARACTERISTIC FAILURE
The illness acting on the brain	Disease course and physiology, read against the timing of the psychiatric change	Establish it and hand back for treatment of the disease, since no psychiatric intervention holds while the process continues	Waiting for a confirmatory test in conditions that have none
The treatment acting on the brain	Exposure and interval, from the drug chart and the procedure record rather than the mental state	Raise the exposure with the treating team, and keep the agent permanently on the differential	Pairing an agent with a syndrome and omitting the interval
A psychiatric disorder in a medically ill person	The patient's own history, baseline and account of themselves, with collateral	Treat it in full as a psychiatric disorder, since this limb is the largest and the least delegable	Accepting that anyone would reasonably feel that way, and stopping there

RULE **Name the limb, then name the evidence that placed you in it.** The follow-up at this interface is one question in different words: what else could be producing this, and what would separate them. An answer that has already given the limb and its evidence has answered it in advance, and bought time to be interesting about the rest.

20.2 The manoeuvres that decide it, in the order they are done

The triage is settled by history far more often than by investigation, and by a small set of manoeuvres that cost nothing and are frequently skipped. They are worth more than any single association here, because they transfer to every section and beyond the chapter.

- **Get a baseline from someone who is not the patient.** A medically unwell patient is observed by staff who never met the well version, so the comparison that would make a change visible does not exist on the ward until a psychiatrist goes and gets it.
- **Read the tempo, and read it for mismatch.** A presentation moving faster than the condition it resembles, continuing to move while adequate treatment is given, or following a slope rather than an episode, is telling you something.
- **Lay the psychiatric change against two timelines.** The medical course with its events dated, and the drug chart read for start dates rather than for the current list. This is the single act that separates the first limb from the second.
- **Ask what was stopped as well as what was started.** An admission removes alcohol, nicotine and whatever else was used at home, and often interrupts the patient's own long-term treatment. That syndrome, on a medical ward, is read as a complication of the illness.
- **Take consciousness and attention before anything else in the mental state.** A fluctuating disturbance of attention and awareness changes the entire answer and moves the problem into territory this chapter names and does not teach. Recognising the shape and handing it on is the examinable act.

TRAP **Settling for understandable.** The commonest error at this interface is not a wrong diagnosis but an abandoned one: the low mood is judged a reasonable response to a serious illness and the assessment stops. Comprehensibility is not a diagnostic exclusion and would not be accepted for any other comorbidity. It costs more here, because the psychiatric state is often what removes the patient from the treatment that would have changed the medical trajectory.

20.3 Telling the systemic conditions apart, which is how they are examined

The autoimmune and inflammatory band is not examined by asking what each condition is. It is examined by asking how you would tell them apart with a patient in front of you, which is why the six conditions sit together on one comparator plate. Two axes organise them: where the discriminating information sits, and what kind of claim you hold when you speak.

CONDITION	WHERE THE DISCRIMINATION SITS	THE SEPARATING SENTENCE	WHAT YOU MAY NOT SAY
Neuropsychiatric lupus	In the reasoning, not in any test	Diagnosis is **attribution**, built from timing against disease activity, absence of a competing explanation, and the character of the syndrome	That any investigation confirms it, or that one manifestation is the commonest
Neurosarcoidosis	Inside the mental state you are already examining	A picture crossing the confusional, psychotic and catatonic layers together, and it can present before the systemic diagnosis exists	That the presentation is common or typical, since a presenting feature is not a frequent one
Primary vasculitis of the central nervous system	In a specialty literature this text could not quote	A declared gap: no self-contained quotation explicitly linking it to a psychiatric manifestation was located	Any named psychiatric syndrome, and equally that it has none, since a failed search is not a negative finding
Antiphospholipid syndrome	Outside the mental state, in the record and the serology	The reported psychiatric range reads like a roster of general adult psychiatry, so the mental state is where the discrimination is not	That the antibody caused the syndrome, since the source reports co-occurrence and no proportion
Coeliac disease	Not in a neuropsychiatric sign at all	A co-occurrence claim against healthy controls, spanning neurodevelopmental, affective, anxiety and eating-disorder categories	That it is disease-specific or causal, since the comparison against another chronic illness was never made
Inflammatory bowel disease	Not in a sign either, but in the course over time	A covariation claim with two arrows and a qualifier: effects run both ways and may influence disease course and psychological health	That psychiatric treatment alters the natural history of the bowel disease

Read the third and fourth columns before the first. Two of the six carry no neuropsychiatric sign of their own, and that asymmetry is the point of putting all six on one plate, because the row not behaving like its neighbours is the row an examiner will use. One condition sits adjacent to the family without belonging to it: the paraneoplastic syndromes are immune-mediated, but their defining feature is a sequence rather than a sign or a claim type, since the syndrome arrives before the malignancy is detected, which is why suspicion, and not the workup, is the psychiatrist's part.

RULE **Name the kind of claim before you name its content.** Reported in is not caused by, an association is not a mechanism, a co-occurrence is not a covariation, and a presenting feature is not a frequent one. This distinction carries more marks than any single disease fact, because almost every systemic association here is reported at the level of co-occurrence. An answer keeping the claim and the clinical move apart reads as trained; one fusing them reads as memorised.

20.4 The treatment limb, compressed to exposure and interval

The second limb has one shape wherever it appears. The decisive evidence is in the drug chart and the procedure record, the question is what was given and when relative to the change, and the detail of what to do belongs elsewhere. What you carry is an agent family paired with a syndrome and with an **interval**, meaning the time between that exposure and the psychiatric change.

EXPOSURE	THE QUESTION IT FORCES	WHERE THE DETAIL IS TAUGHT
Calcineurin inhibitors after transplantation	What else was given, in what order, and when relative to the event, since one labelled signal is conditional on a combination	Label level in the transplant section; prescribing and interactions in the psychopharmacology notes
Corticosteroids	Which agent, over what course, and what was the earlier psychiatric history	Routed out of the chapter to the endocrine and metabolic notes, so there is exactly one place to look
Cytotoxic treatment	Did the difficulty predate the first cycle, does it fluctuate within a day, is it evolving, and has it an objective correlate	The cancer-drug section, which teaches the construct rather than the patient's word for it
Interferon-alpha	Prevention before a defined course or treatment after the syndrome declares itself, and in which population	The cancer-drug section; interferon-beta belongs to the neurology and psychiatry interface notes and is never pooled with it
Anaesthesia and surgery	Was anything measured before the operation, since a change is a relation between two observations	The perioperative section for the construct; the delirium notes for the acute confusional presentation
What the admission removed	What was stopped, at what hour, and whether the timing fits	Here, in the triage, since no other chapter owns a removal as an exposure

TRAP **Naming the agent and the syndrome and stopping.** That pairing is the least discriminating half of the topic and every candidate has it. The interval tells a clinician when to look and which explanations compete at that moment, and the population tells them whether the finding applies to this patient. A finding stated without its population is not a shorter answer, it is a weaker and different one.

20.5 The boundaries, which are a competence and not an apology

Liaison work is routing, so a chapter that practises routing is teaching the clinical skill rather than excusing an omission. A referral arrives from another team, written by someone unsure what to ask, and the first act is deciding whose question it is: the treating team's, the psychiatrist's, or a third specialty's. A resident who has practised deciding where a construct belongs has been practising the job; one who memorised a thin paragraph about each has not, and the paragraph fails the follow-up. Stating the boundary aloud also makes an answer checkable: whether a chapter mentioned a topic is a test it can pass while leaving its reader unable to answer anything, whereas whether the reader can name the owner has a right answer.

THE QUESTION IN FRONT OF YOU	WHO OWNS IT
Which agent, at what dose, adjusted how for renal, hepatic, cardiac or transplant disease	The psychopharmacology notes. This chapter teaches everything preceding a prescription and adds no second version of it.
Is this an acute confusional state, and how is it assessed and managed	The stroke, traumatic brain injury, delirium and focal brain syndromes notes, including every bedside instrument, rating scale and drug choice.
Does this patient have capacity, and what is the statutory test	The mental health legislation notes. This chapter says only that capacity is decision-specific and time-specific.
The general immune-to-brain model, and the workup of an autoimmune encephalitis	The neurology and psychiatry interface notes, which also own interferon-beta. This chapter owns the disease-specific presentation and the decision to suspect.
Depression and anxiety after infarction, critical-illness sequelae, dialysis, depression and demoralisation in cancer, palliative care psychiatry	The consultation-liaison, geriatric and transcultural psychiatry notes. These are defined by a setting and a referral pathway, not by an organ.
Psychiatric disturbance following a corticosteroid	The endocrine and metabolic psychiatry notes. This row follows a drug rather than a setting, and follows it out of the chapter.
A possible assault, and the pathway that follows it	The forensic and medicolegal material. Raised in the burns section as an antecedent question and handed on at once.

RULE **Route in the answer, out loud, before you are asked to.** A cross-reference is an address and not an abstract, so it will not survive being revised from, and the failure mode is a candidate who knows a construct matters but cannot say what it does. Say what you are handing on, who holds it, and what you keep. That costs seconds and demonstrates the judgement this interface examines.

20.6 The declared gap, and the sentence that carries it

This chapter declares more gaps than most, deliberately, and the declarations are examinable content rather than housekeeping. A candidate who can say what is not established, and on whose authority a statement would rest, answers better than one who fills the space, because a filled space is unverifiable and frequently wrong. A **declared gap** has three parts, each load-bearing.

State what was sought, with its boundaries: a quotable statement of a particular kind, linking a named condition to a psychiatric manifestation, in a source of stated standing, rather than a vague admission of uncertainty. Name whose literature would hold it, since the specialty following these patients owns the phenotype and deferring to it by name is a different act from not knowing. Then say what you do anyway, because a missing figure never changes the clinical move: prevalence tells you how hard to look, never whether this patient has the syndrome.

The declared gaps are worth revising as a set, since a stem often lands on one. The psychiatric phenotype of primary vasculitis of the central nervous system was not obtained, nor was a specialty statement binding psychiatric management to chronic respiratory failure. No prognostic figure is attached to anxiety after infarction, and no usable recommendation was opened on the desire for hastened death. The epidemiological specifics ordinarily quoted about the stress-precipitated cardiomyopathy are absent, and that section teaches what such a figure would have to measure. In lupus the antibody associations and every frequency statement are absent, because without a diagnostic gold standard each study defines its own numerator. In burns the prevalence figures and the work-reintegration comparison were not obtained, while the Indian work on self-immolation was. And no mechanism is asserted anywhere in the gut-brain material, because neither finding it holds identifies a pathway.

TRAP **Turning a missing number into a negative finding.** The candidate who cannot recall a figure says the association is weak or the condition is rare. That is a categorical claim built on a failure of memory, and it is unsafe where being wrong costs a life. Absence of a located quotation is a fact about a text's sources, never about the disease.

20.7 Answer structures, and the order the reasoning goes in

Almost every question at this interface is one of four stems, and each has an opening move that buys the rest of the answer. Learn the openings, because the first two sentences are where an examiner decides whether you are reasoning or reciting.

- **Discuss the psychiatric aspects of a named medical condition.** Open with the triage applied to that condition, naming which limb it mostly trains and why. Do not open with a list of syndromes that could have been written about any organ.
- **When would you suspect an occult or systemic cause.** Open with the readings that raise suspicion: tempo, the company the syndrome keeps, and whether a general history was taken at all. Then say what you would do, and route the investigation to its owner.
- **How would you tell these conditions apart.** Open by naming them separately and refusing the single undifferentiated paragraph about inflammation reaching the brain. Give each its discriminating feature, and mark which parts rest on a source and which on clinical reasoning.
- **A patient with a known systemic diagnosis presents with a psychiatric syndrome.** Open with the exclusions, in the order infection, metabolic disturbance, the treatment itself, and primary psychiatric illness, because those are commoner and more reversible than the manifestation being reached for.

The body then runs in the same order whichever stem opened it. Characterise the syndrome in the patient's own words. Place it against the three limbs and say what put it there. Give the discriminating feature separating it from its nearest neighbour. State each finding with its population and observation window attached. Name every construct you hand on, with its owner. Declare the gap where the chapter has one, in its three parts. Then say what you would do on the ward tomorrow, which is what most answers run out of time for and what a marker is waiting on.

Two sentences are worth rehearsing aloud, because candidates who have never said them tend to bluff. The gap sentence: that figure is held in the specialty's literature, I cannot state it from a source, and it does not change what I do with this patient. The attribution sentence: this is my judgement rather than a finding, and here is what would change it.

Both sentences are built from the same four fields, and so is every finding this page asks you to state. The comparator grid asked what kind of claim each row holds, the exposure table asked when and in whom, and the gap declaration asked what is not established and whose literature would hold it. Four words carry all of it.

MNEMONIC

Claim, Population, Window, Gap

Claim: say what grade of statement you are holding before you say what it contains, since the grades separated in the comparator rule above are not interchangeable and the weakest of them is the one most often reported here. **Population**: say in whom it was observed, because leaving that out does not shorten the answer, it changes what the answer means. **Window**: give the period over which it was seen, and on the treatment limb the gap between the exposure and the change, which is the half of that pairing most candidates omit. **Gap**: where nothing can be stated, declare it in the three parts set out above. Asked in that order the four survive a condition this chapter never taught, which is why they outlast any single row in the tables.

20.8 If there is an hour left

Revise in this order and stop when the hour is gone rather than when the chapter is finished. The order puts first what an unfamiliar stem is most likely to need.

- **The triage and its evidence table.** It is the answer structure for every question here, it works on constructs you have never met, and it is the shortest thing on this page.
- **The manoeuvres.** Baseline, tempo, the two timelines, what was stopped, and consciousness first.
- **The comparator grid for the six systemic conditions.** Work across each row, ask what makes that condition itself rather than the one above it, and check what kind of claim the row is.
- **The owner list.** Prescribing, delirium, capacity, the immune-brain model with the encephalitis workup, the setting-defined liaison constructs, and the corticosteroid.
- **The gap sentence in its three parts,** and the specific gaps this chapter declares, since a stem landing on one is a stem you answer well while others invent.

- **The exposure and interval table,** read for the second column, because agent to syndrome is the half everybody has and the interval is the half that separates.

What can safely be left is every list this chapter did not build an answer around, and every construct owned elsewhere. Reading a routed topic here in the last hour costs the material only this chapter carries: systemic and inflammatory disease presenting psychiatrically before anybody knows it is present.

If only one card leaves the hour with you, make it this one: four questions to ask of any finding before you say it aloud.

Claim	Population	Window	Gap
WHAT KIND, BEFORE WHAT CONTENT	IN WHOM IT WAS OBSERVED	OVER WHAT PERIOD, AND HOW LONG AFTER	NOT ESTABLISHED, AND WHOSE LITERATURE HOLDS IT

IN INDIA

No India-specific figure is held for this consolidation and none has been invented. The Indian point is about information rather than epidemiology. In most Indian general hospitals the psychiatrist reaches the medical bedside on request, usually late, to a referral question written by a team under pressure, and often without the previous notes, the collateral history or the drug chart's start dates. Every manoeuvre above depends on information the referring team holds and the letter rarely contains, so the skill is asking for it explicitly and early rather than working from the mental state alone. Distance and cost then decide whether the plan is followed at all.

GOOD TO REMEMBER

Three questions at the bedside, one afterwards, and one sentence held in reserve. Is the disease doing this? Is the treatment doing this? Is this a psychiatric disorder in a person who is unwell? Then, before writing: whose chapter answers the part I am handing on? And in reserve, for the moment a figure is wanted and cannot be sourced: that number is held in the specialty's literature, I cannot state it, and here is what I would do regardless.

The chapter installs a triage, a habit of naming owners, and a way of declaring what is not established. Everything else in it is worked examples, and a candidate who holds those three can answer on a condition this chapter never mentioned.

Index

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

A

access 11
access pathway 11
activity returns 230
acute stress disorder 318
adjustment 111
administrative counts 181
adolescent-friendly service 169
age standardisation 89
allocation 61
antiphospholipid syndrome 240
anxiety sensitivity 255
appearance concern 319
assertive follow-up 120
associative stigma 92
attendance management 176
attribution 14
avoidance 250

B

bidirectional referral 130
bidirectionality 240
bounded 80

C

cancer-related cognitive impairment 284
capacity-based return 176
care manager 155
carer organisations 131
case definition 49
caseload-based 155
catalogue answering 233
catecholamine surge 263
chemobrain 283
ciclosporin 270
clinic-attender sample 214
co-occurrence claim 312
cognitive and affective items 248
collaborative care 137
company 21
competency 82
competitive employment 110
continued use 200
continuity of care 42
contraction band necrosis 264
cost of disclosure 175
cost of illness 186
cost-effectiveness 154
covariation claim 312
coverage 13
crude rates 89
cumulative activity figure 150

D

day care 32
decay 100
decision-specific and time-specific 241
declared gap 20
decriminalisation 28
deinstitutionalisation 108
delivery-side distance 228
denominator 13
digital divide 198
disability 28
discrimination 91
disfigurement 319
displacement 28
dissociation 318
district mental health programme 20
dose 40

E

economic evaluation 186
efficacy belief 103
endpoint selection 100
environments people live through 203
escalation route 141
explanatory-model assessment 213

F

fidelity 110
first-response knowledge 93
fit 14
flow 34
functional coverage 15

G

gatekeeper 68
global mental health movement 203

graded 115
gradient 71
guided or blended intervention 198
gut-brain axis 309

H

halfway home 110
handover 12
hap 11
harm reduction 216
healer training 205
high-risk group 85
hub-and-spoke model 148
hubs 148

I

india's central occupational-safety statute 162
india's national mental health programme 189
indicated 73
indigenous resources 129
individual placement and support 110
inflammatory limb 302
insecticides act, 1968 80
institutionalism 121
insurance parity 44
integration 14
intentional non-adherence 257
interferon alfa 286
interferon-alpha 283
interferon-beta 242
internal migration 205
internally displaced people 179
interpreter 162
intersectoral action 207
interval 29
interval between onset and effective treatment 211
irdai 44

J

job tenure 118

L

layered response 161
limited awareness, sociocultural beliefs, values and stigma 56
locally meaningful idioms of distress 205

M

measurement proxy 231
mechanistic claim 312
media guideline 78
mental health literacy 59
mental-status examination 271
mhgap 137
microvascular dysfunction 264
model of care 130
multisectoral 73
myocardial stunning 264

N

national crime records bureau, the ncrb 85
national health mission 59
national mental health programme 17
national mental health survey 36
national suicide prevention strategy 66
need-side distance 228
need-to-service distance 227
neurosarcoidosis 240
norms of journalistic conduct, 2020 80
npsle 298

O

occupational health function 177
onconeural 290
onconeural antibody 292
order in time 325

P

paraneoplastic 239
participation 93
pathway to care 199
perioperative neurocognitive disorders 278
perspective 186
phantom shock 251
platform intervention 167
police-recorded suicide cases 85
population lever 78
population register 155
population-level intervention 225
portable record 180
posterior reversible encephalopathy syndrome 272
preparedness 164
primary central nervous system vasculitis 301
primary prevention 184
protracted situations 179
psychological first aid 159
psychosocial hazards 174
public stigma 91
public-health implementation 11
pulmonary rehabilitation 256

R

rapid needs assessment 160
reactance 99
recognition 11
recognition capacity 211
recovery 59
reference year 13
referral word 330
release 18
remote consultation 34
requires trust, rapport-building and open dialogue 205
responder care 165
route knowledge 103
routed construct 239
routing 14

S

safeguarding floor 217
safety escalation 201
sanctioned coverage 15
secondary prevention 185
selective 73
self-guided intervention 198
self-stigma 91
sequencing rather than selection 275
service architecture 11
service variable 121
shared-device problem 199
shock-related anxiety 250
social determinant 34
special population 178
spokes 148
statutory entitlement 27
stepped care 135
stock 20
structural adaptation 144
structural measures 97
structural stigma 92
suicide surveillance 66
supervision 14
supported employment 32
surface adaptation 144
symptom-rhythm concordance 250
symptomatic limb 302

T

task-sharing 137
task-shifting 19
tele-manas 34
tempo 166
tertiary prevention 184
thrombotic limb 302
to enhance human resource in mental health sub-specialties 18
total body surface area 321
trained lay counsellor 137
transfer 19
transinstitutionalisation 110
transition 63
transportability 117
treatment gap 14
triage 39

U

undercount 87
unintentional non-adherence 257
unique users 36
universal 73
uptake 28
utilisation 28

V

valid denominator 13
voluntary sector 20

W

what a service has 27
what a service owes 27

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

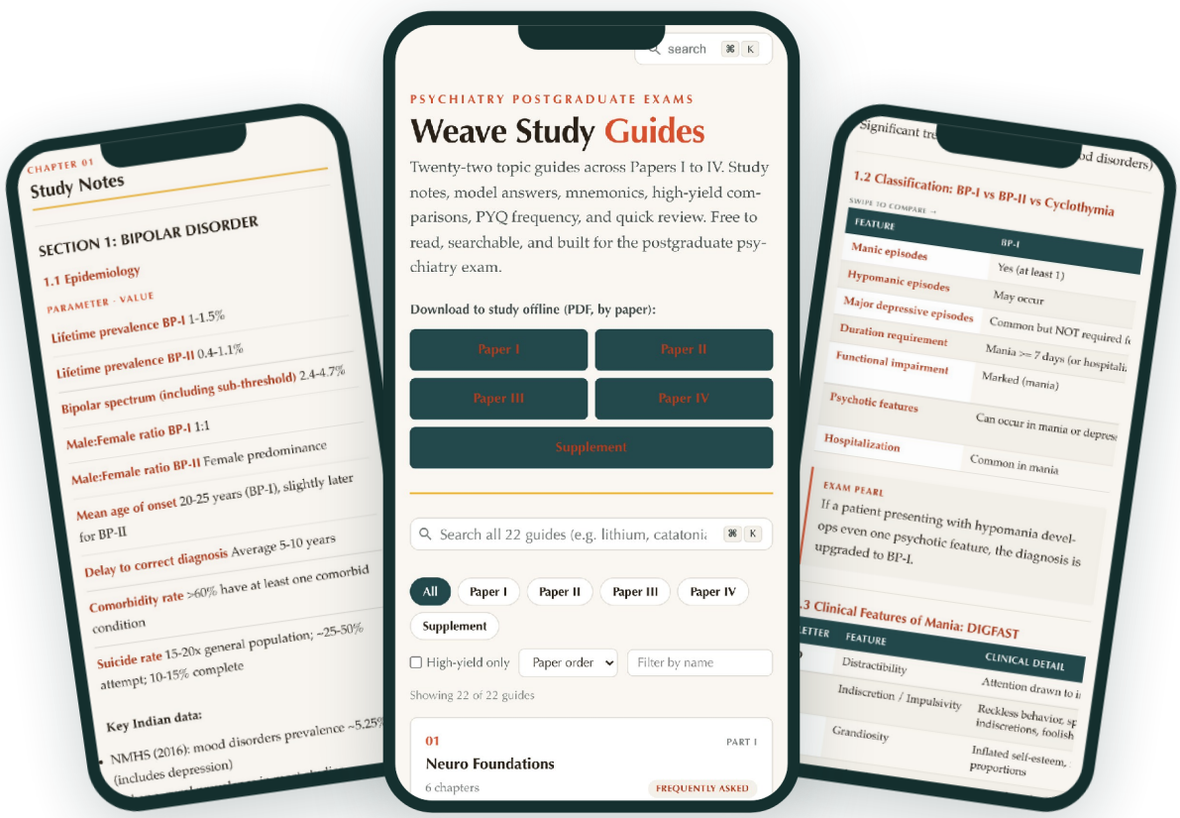
weave.clinic/library

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

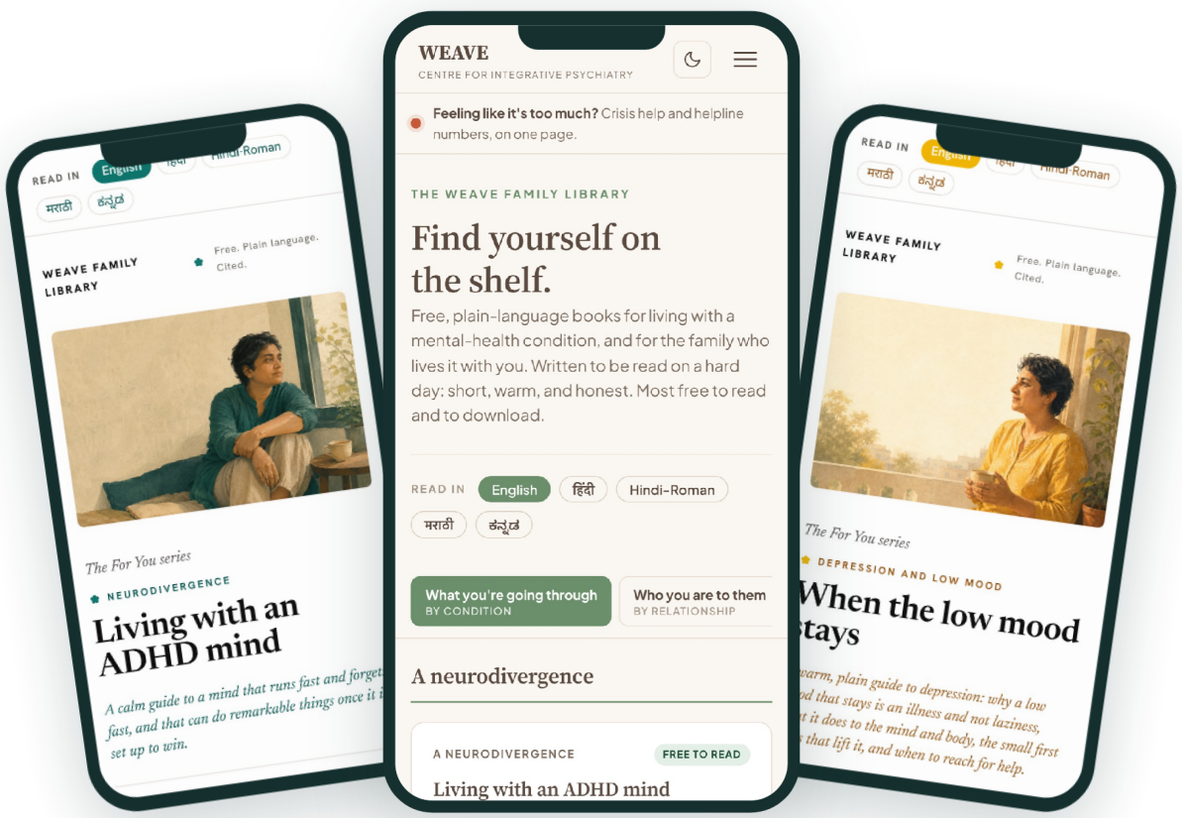
weave.clinic/study

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 SEVEN

Volume Seven of the fourth paper closes the set: community psychiatry and public health, the burden of disorder, the treatment gap, the national programmes meant to close it and prevention that is more than advice; then psychiatry at the medical interface, the heart, the lungs, the failing organ, the transplant and the cancer. Two complete sets of notes, written to be understood the first time and recalled every time after, with a colour memory code, a clickable contents and a full index.

Built to understand. Kept to revise.

All the very best for your exam.

WEAVE

weave.clinic · wilfred@weave.clinic · niharika@weave.clinic