

FLOW HIJACKED

BEFORE TOMORROW HAPPENS

Anxiety, Prediction, and the Shrinking Future

A Flow Hijacked Monograph

Tomer Avraham, PhD

September 2026

Contents

Prologue	8
PART I: THE FUTURE ARRIVES EARLY	11
Chapter 1 - Before Anything Happens	12
Chapter 2 - Fear, Anxiety, and the Difference Between Now and Next	19
Chapter 3 - When Prediction Becomes Policy	26
PART II: LEARNING DANGER	32
Chapter 4 - How Fear Learns	33
Chapter 5 - Safety Is Learned Too	39
Chapter 6 - Memory, Context, and the Return of Fear	45
PART III: WHEN SAFETY SHRINKS THE WORLD	53
Chapter 7: When One Danger Becomes Many	54
Chapter 8: Avoidance, the Solution That Works Too Well	59
Chapter 9: Worry, Reassurance, and the Search for Certainty	64
Chapter 10: The Body Before the Explanation	69
Chapter 11: Interoception, When Sensation Becomes Evidence	73
Chapter 12: Panic, Catastrophe From Within	77
PART IV: THE PREDICTIVE FUTURE	81
Chapter 13: The Threat-Prediction Machine	82
Chapter 14: The Future as a Place of Danger	86
Chapter 15: Uncertainty Is Not Neutral	90
Chapter 16: The Shapes of Uncertainty	94
Chapter 17 - Prediction, Priors, and the Weight of Evidence	98
Chapter 18 - Bayes and the Committee of Possible Causes	106
Chapter 19 - Reinforcement Learning Under Threat	114

Chapter 20 - The Price of Not Exploring	123
PART V: MANY FACES OF AN ANXIOUS FUTURE	130
Chapter 21 - The Future With No Closing Time	131
Chapter 22 - Threat From Within, Threat From Other Minds	137
Chapter 23 - Narrow Fears, Overlapping Worlds	143
Chapter 24 - Ancient Alarm, Twenty-First-Century Bandwidth	150
PART VI: THE FLOW HIJACKED MODEL	157
Chapter 25 - The Anticipatory Attractor	158
Chapter 26 - How the Future Becomes Smaller	164
Chapter 27 - Anxious Hysteresis	171
PART VII: MAKING THE FUTURE REACHABLE AGAIN	178
Chapter 28 - Treatment as New Learning	179
Chapter 29 - Acceptance, Metacognition and the Freedom Not to Obey	185
Chapter 30 - Changing the Conditions of Change	190
Chapter 31 - When Treatment Does Not Work as Expected	196
Chapter 32 - Making the Future Reachable Again	202
Epilogue - Before Certainty	207
Evidence Notes and References	209

Why I Wrote This Book

There are subjects one approaches because they are intellectually interesting, and subjects one approaches because they have already entered the room. For me, the shrinking of the future belongs to the second kind.

I did not arrive at *Flow Hijacked* from the safe distance of a seminar table. I came to it through years in which prediction, bodily state, addiction, depression, shame, treatment, and the long reconstruction of agency were not separate topics. They were parts of one life. I spent a year and a half in a therapeutic community. During the first year I lived without a phone. Later came workdays outside the community, the awkward return of ordinary responsibility, sport, meditation, writing, relationships, and time with my children. Recovery did not arrive as a grand revelation. It arrived as a field of possible actions that had to become reachable again.

That experience does not prove a scientific model. A life cannot substitute for a controlled comparison, and suffering does not confer methodological authority. But experience can expose a question that academic categories sometimes divide too quickly. What happens when the future is not merely imagined but begins to organize the present? What happens when a local act of protection is rewarded so efficiently that the larger world grows smaller? Why can a person understand a danger intellectually and still find that the body, attention, and available actions have already voted another way?

Anxiety gives those questions a particularly clear form. It is often described as excessive fear, distorted thinking, or a problem of uncertainty. Each description catches something. None, by itself, captures the moving system. A forecast changes the body. The changed body alters what seems plausible. Attention selects evidence. Action changes which evidence can appear. Relief teaches. Memory reconstructs. Other people regulate, reassure, alarm, and sometimes lend us a nervous system sturdy enough to explore. The environment

decides which futures exist in practice, not only which futures can be imagined.

This book follows that movement from familiar moments into learning theory, interoception, predictive processing, reinforcement learning, development, treatment, and dynamical systems. It does not ask the reader to admire complexity. It uses complexity only when it earns its place by explaining something that simpler language leaves behind. The aim is not to turn a person into an equation. It is to stop mistaking a person for a symptom score.

I am the creator of Flow Hijacked and an independent researcher. I am not writing as a clinician, and this book is not an invitation to diagnose oneself or alter treatment without qualified care. It is a non-profit educational work: an attempt to build a bridge between rigorous evidence and the felt architecture of a life. Wherever the evidence is strong, the language should be strong enough to say so. Wherever the evidence is mixed, preliminary, indirect, or absent, the language should become correspondingly modest. The original Flow Hijacked concepts introduced later in the book are hypotheses. Their dignity depends not on sounding impressive, but on being precise enough to fail.

The personal thread in these pages is therefore neither confession nor proof. It is a reminder of the stakes. A future can become small long before it disappears. It can remain richly populated in imagination while only one or two actions feel behaviourally available. Conversely, recovery can begin before confidence arrives. A phone call can be made while doubt remains. A room can be entered before the body offers permission. A parent can be present for a child on a day that still contains fear. These are not cinematic victories. They are the quiet reopening of routes.

I have tried to write for several readers at once without treating any of them as a simplified version of another: the person who recognizes the experience from the inside, the loved one trying to understand it, the therapist thinking about process, and the researcher asking whether the proposed mechanisms can be measured. The book moves

in layers because human beings do. A kitchen table, a racing heart, a Bayesian prior, a family accommodation pattern, and a hysteresis loop can belong to the same inquiry without pretending to be the same kind of fact.

The promise here is not a harmless future. No honest book can offer one. The promise is a more navigable account of how futures acquire weight, how protection becomes policy, how corrective evidence is starved, and how a larger field of action may be rebuilt. The future will remain uncertain. The important question is whether uncertainty must remain its only gatekeeper.

A Note on Evidence and Safety

This edition separates three levels of statement. An **empirical result** reports what a particular study or synthesis found within its population, method, and limitations. A **broader interpretation** connects results across literatures while acknowledging that the connection was not directly tested as one mechanism. A **Flow Hijacked synthesis or hypothesis** is an original proposal. It is not presented as an established biological entity or a clinically validated instrument.

Compact evidence markers such as ^[LE-03] point to full source entries at the end of the book. They are not decorative citations. Each marker belongs to a claim-to-source ledger that records what the source can support and what it cannot. Group averages are not individual destinies. Brain-imaging associations are not diagnoses. A short-term treatment response is not the same as durable recovery. Animal learning experiments can illuminate principles without becoming direct explanations of a particular human life.

This book is educational. New or severe chest pain, fainting, substantial breathing difficulty, a new neurological sign, confusion, suspected intoxication or withdrawal, or another medically concerning change should not be assumed to be anxiety. Medication changes, including discontinuation, require qualified medical guidance. In immediate danger, seek urgent local help.

Prologue

The Night Before

At 10:42 p.m., the phone is face down on the kitchen table.

It has not rung. No message has arrived. The room contains a half-finished cup of tea, the soft motor of a refrigerator, a chair pulled slightly away from the table, and one human being waiting for a conversation scheduled for the morning. The conversation may bring ordinary news. It may bring difficult news. At this hour, both futures are still available.

Only one of them has begun furnishing the room.

The tea has cooled because the person keeps forgetting to drink it. A sentence spoken three days earlier is retrieved and examined, then retrieved again with a different emphasis. The jaw is tight. The shoulders have risen. Every few minutes the phone is turned over, although a silent screen cannot disclose tomorrow. A search begins online and is abandoned, then resumed with a narrower phrase. The results provide information, but information is not what the system seems to be asking for. It is asking for an end to uncertainty.

Nothing has happened in the ordinary sense. No verdict has been delivered. No loss has occurred. No door has closed. Yet the event has already acquired physiological weight. It has altered breathing, appetite, attention, memory, and the imaginable shape of the next day. Tomorrow has crossed the calendar boundary early.

This is one of anxiety's strangest powers. It can charge the present for an event that does not yet exist.

The power is not foreign to intelligence. We live by reaching forward. We leave home before a train departs. We save for winters that have not arrived. We rehearse difficult conversations, insure houses, pack medicine, teach children to cross roads, and plant trees whose shade belongs to people not yet born. A creature unable to let the future influence the present would not be admirably spontaneous. It would be helplessly late.

The question, then, cannot be how to stop anticipating. It must be how anticipation is governed.

Back at the table, the mind continues its work. It produces a bad outcome, then asks what could prevent it. It produces a second bad outcome that survives the first precaution. It examines a neutral detail and discovers that the detail is compatible with danger. That compatibility feels important. Soon, mere possibility begins to function as an instruction.

Check once more.

Do not make plans for after the call.

Stay awake in case preparation is still possible.

The instructions are not absurd. Each offers a small form of protection. Checking may reveal something. Cancelling tomorrow's plans may prevent embarrassment. Staying vigilant may avoid being surprised. The problem is not that these actions make no local sense. The problem is the kind of world they gradually construct: a world in which safety depends on continued vigilance, and in which the feared future is rarely allowed to meet an unprotected reality.

At 2:13 a.m., the phone remains silent. That silence could mean nothing. It could even be mildly reassuring. But silence has no stable meaning when the mind is required to prove that the future is safe. It becomes an empty space into which another simulation can be placed.

Morning will eventually settle the immediate question. Perhaps the call will be difficult. Perhaps it will be routine. Either outcome matters. But the deeper subject of this book lies in what happened before the answer: how a nervous system built to help a person reach the future can begin reducing the number of futures the person is willing or able to enter.

Anxiety is often described as too much fear. Sometimes it is. But intensity alone misses the architecture. Anxiety can change which memories are searched, which sensations become evidence, which actions feel available, which uncertainties are tolerated, and which parts of life are postponed until a guarantee arrives. It can make a

single possible tomorrow more influential than the many actual things occurring today.

This book follows that change.

It begins before anything happens.

PART I

THE FUTURE ARRIVES EARLY

How possibility begins charging the present.

Chapter 1 - Before Anything Happens

Before Anything Happens

The future sends invoices early

The dentist appointment is on Thursday, but the jaw tightens on Monday. A performance review lasts forty minutes, yet it can occupy the week before it. A train journey that will probably be uneventful can borrow sleep from the night before departure. We speak casually of events “weighing on” us, as if Thursday had mass and Monday were a surface on which it could rest.

In a biological sense, the metaphor is not far wrong. Anticipated events can recruit present resources. Attention is redirected. Muscles prepare. The heart changes pace. Memory searches for relevant precedents. Plans are revised. The future has no physical access to the present, but a model of the future does.

Foresight is therefore not an ornamental capacity added to an otherwise reactive animal. It is part of how action is organized. If thirst were noticed only after severe dehydration, if a fast-moving vehicle mattered only after impact, or if winter prompted preparation only when food was already scarce, regulation would fail. Living systems survive partly by paying some costs in advance.

Anxiety begins with this respectable assignment: look ahead; detect costly possibilities; prepare before preparation becomes useless.

The difficulty is calibration. Which futures deserve an advance payment? How large should it be? What counts as enough preparation? And what happens when the system keeps paying after every useful action has already been taken?

Consider a small message from a manager: “Can we talk tomorrow?” It contains almost no information about whether the conversation will be pleasant, unpleasant, or forgettable. That absence of information gives imagination room. One person reads it, checks

the calendar, and returns to dinner. Another begins an internal audit. Was the report late? Did the manager's voice sound cool yesterday? Why was a senior colleague copied on that email? A seven-word message becomes a search instruction.

The difference need not be explained by one person being rational and the other irrational. They may have different histories, current stresses, beliefs about authority, bodily states, or estimates of the cost of losing a job. The point is what follows. A forecast reorganizes sampling. Memory is searched for mistakes rather than praise. Ambiguous expressions are inspected for warning. Sleep becomes lighter. By morning, the body itself feels like additional evidence that the situation must be serious.

Predictions do not merely wait in the mind for reality to grade them. They can help assemble the evidence environment in which they will later be judged.

Preparation and the missing stopping rule

Useful preparation usually contains a stopping rule. A bag is packed. A route is checked. A talk is rehearsed twice. An appointment is made. More preparation eventually adds little, and attention returns to the rest of life.

Anxiety can move the stopping rule without announcing that it has done so. The practical question, "Have I prepared adequately?" becomes the impossible question, "Have I made this future safe?"

The difference is decisive. Adequate preparation is often finite. Absolute safety in an open future is not.

Suppose a person is flying the next morning. Checking the departure time, bringing a passport, and allowing enough time to reach the airport are bounded tasks. Once completed, they do not answer whether the aircraft will certainly land. Nothing can. If the nervous system treats that unanswered question as unfinished work, preparation continues after the work has lost its object. Weather maps are refreshed. Mechanical sounds are researched. Reviews of airlines are read as if customer ratings could close probability itself.

The same migration occurs in other domains. Responsible attention to a symptom becomes repeated checking for a guarantee that no future illness can appear. Care for a relationship becomes monitoring for proof that another person's feelings will never change. Conscientious work becomes an attempt to prevent every possible criticism. Moral reflection becomes a demand to know with certainty that one has never caused harm.

In each case, anxiety attaches itself to something valuable: health, love, competence, responsibility. This is why the instruction to “stop caring” is both cruel and conceptually confused. The problem is not the value. It is a control strategy that asks the future to become closed before life may proceed.

A prospective control problem

To call anxiety a prospective control problem is not to compare a suffering person to a thermostat. It is to notice a shared question: how should current action change in light of a difference between an expected state and an acceptable one?

Human reference conditions are exceptionally rich. We regulate not only temperature, energy, pain, and physical distance from harm. We care about belonging, dignity, competence, guilt, abandonment, the safety of people we love, and the kind of person we may become. Symbolic events can therefore recruit flesh. A possible humiliation can raise a pulse. An imagined failure can remove appetite. A sentence expected tomorrow can tighten the chest tonight.

“It is only in your head” fails because representations in the nervous system are not weightless spectators. They participate in the regulation of the body and in the selection of action.

Even so, there is no single proven “anxiety circuit” issuing commands from a control room. Imaging studies of induced and clinical anxiety point to distributed and partly overlapping patterns involving systems concerned with salience, bodily state, action, and control. Those group-level patterns vary across tasks, diagnoses, and

individuals; they are not scans from which one can read a person's private future. ^{[NB-01][NB-02]}

That is an empirical result with a boundary: anxiety engages distributed systems, but distributed activation does not by itself explain why a particular person worries about a particular event.

The interpretation offered here is functional. Anticipation changes present regulation because expected futures are treated as relevant to current action. This interpretation joins evidence from several literatures; no experiment has established one master prospective-control mechanism called anxiety.

The Flow Hijacked synthesis goes one step further. It proposes that the clinically important shift is often not simply an increase in unpleasant feeling but a change in the reachable field of action. A threatening future gains enough weight to make some paths difficult to enter at all. This is an organizing hypothesis, not an established biological entity. It will need measurable definitions and the possibility of failure before it deserves more than provisional trust.

Busy without moving

An anxious day can feel exhausting even when little happens externally. The apparent stillness is misleading. The system may be generating scenarios, comparing precautions, monitoring the body, reviewing conversations, suppressing distractions, checking for news, and preparing escape routes. It is running a costly simulation department.

The department has a peculiar accounting system. A feared possibility can demand resources merely by remaining possible. A competing possibility—an ordinary meal, a friendly conversation, an absorbing task—may need to prove that it is safe enough to deserve attention. The threshold is unequal.

This helps explain why advice to “focus on something else” often lands badly. Attention is not simply refusing instruction. It has been recruited by a model that assigns unusually high cost to missing a

threat. If the estimated consequence of overlooking one clue is catastrophic, returning to a novel or a conversation can feel negligent.

The system may also mistake activity for progress. Another search creates the feeling of doing something. Another rehearsal resembles preparation. Another scan of memory appears to reduce the chance of surprise. Yet without a stopping rule, activity can continue while uncertainty remains unchanged.

There is a useful distinction here between solving and circling. Solving transforms the problem: a call is made, information is obtained, a practical action is completed. Circling repeats control attempts that cannot resolve the open variable. From inside the loop, however, the two can feel almost identical. Both involve effort. Both are motivated by care. Only one changes the world in a way that can complete the task.

The body votes on plausibility

Imagine the manager's message again after a poor night's sleep. The same words are on the screen, but they do not arrive in the same organism. Fatigue makes concentration harder. The heart is faster after too much coffee. The chest feels tight. These changes do not prove that the meeting is dangerous, but they can alter what feels believable.

This is not because the body is a faulty witness that should always be ignored. Bodily signals are indispensable. They report exertion, illness, hunger, pain, arousal, and countless changes that matter. The difficulty is inference. A racing heart tells us that the heart is racing. It does not arrive labelled "because the manager intends to fire you," "because you slept four hours," or "because this is a medical emergency." Meaning must be assigned.

That assignment is influenced by the forecast already in play. If the leading model is danger, arousal is likely to be read as confirmation. The confirmation strengthens attention to danger, which can increase arousal further. Soon the loop has produced genuine evidence of distress without producing evidence about the manager's intention.

The distinction is subtle but humane. The distress is not imaginary. The causal conclusion remains uncertain.

What anxiety is protecting

Many accounts of anxiety begin by asking what has malfunctioned. A better opening question is often what the system is trying to protect.

Perhaps it is protecting a body from illness, a bond from rupture, a child from harm, a reputation from humiliation, or a self-concept from moral failure. The answer matters because protection is organized around values and histories, not just probabilities. Two people can assign the same likelihood to an event and experience it differently because its meaning and expected cost are different.

This does not make every precaution adaptive. It explains why precautions can remain compelling after an observer has declared the probability low. Low probability multiplied by intolerable meaning can still dominate a decision.

The appropriate scientific task is therefore not to sneer at distorted thinking, nor to baptize every alarm as wisdom. It is to examine how estimates of probability, cost, controllability, coping, and uncertainty interact; how those estimates change the body; and how subsequent action changes the evidence available for learning.

Three levels of claim

The distinction among evidence, interpretation, and synthesis will matter throughout this book.

Empirical result: anticipated and induced threat recruits distributed brain and bodily systems, and the patterns observed across studies are heterogeneous rather than reducible to one universal locus. ^{[NB-01][NB-02][NB-07]}

Broader scientific interpretation: anxiety can be understood as prospective regulation in which possible future costs alter present attention, physiology, and action. This interpretation is consistent with the evidence, but it is not a single measured variable shared by every study.

Flow Hijacked synthesis: anxiety may become self-restricting when protective policies reduce the range of actions and evidence through which threatening predictions could be revised. That proposal will not be treated as a diagnosis, a discovered brain module, or a substitute for older science. It is useful only if later chapters can turn it into risky, testable questions.

For now, one conclusion is enough. Before an event happens, a representation of it can change the state from which the event will eventually be met.

The future arrives early. In the next chapter, we must ask what distinguishes the future danger of anxiety from the danger that is already here.

Chapter 2 - Fear, Anxiety, and the Difference Between Now and Next

Fear, Anxiety, and the Difference Between Now and Next

A tiger, a timetable, and a useful imperfection

Fear is usually introduced with a predator. A tiger steps onto the path; the organism mobilizes. Psychology has required a great deal of unpaid teaching from tigers.

The example survives because immediacy is easy to see. There is an animal, a distance, a direction of movement, and an urgent decision. Anxiety is more difficult to point at. Its object can be absent. Tuesday's journey can alter Monday's body. A test result that has not been produced can govern the hours before it. The feared thing may not be an object at all, but a transition: calm becomes panic, belonging becomes rejection, health becomes illness, control becomes exposure.

It is tempting to draw a bright line: fear belongs to present danger; anxiety belongs to future danger. Biology declines the invitation. The two share defensive physiology, attention, learning, memory, and action systems. Immediate and uncertain threats often blend. A person waiting outside an examination room is facing something more proximal than a person worrying six months earlier, yet the result remains unknown. A person walking home during a period of credible external threat may be both responding to real conditions and simulating what could happen next.

The distinction is therefore a gradient, not a border. It is useful because time, distance, and certainty change the control problem. In an experimental task involving a simulated predator, more distant threat was associated with greater recruitment of anterior regions, while threat that drew near was associated more strongly with midbrain and periaqueductal systems. This was a small laboratory study, not a map of daily fear, but it supports the idea that defensive organization changes with threat imminence. ^[NB-05] Studies of temporally uncertain anticipation likewise implicate a distributed set

of cortical, subcortical, and brainstem regions rather than a neat anatomical switch labelled “anxiety.” ^[NB-07]

Empirical result: distance, timing, and uncertainty can alter patterns of defensive responding, and those patterns are distributed. **Interpretation:** fear and anxiety are overlapping modes of prospective defence whose balance shifts with the structure of threat. The first statement is more secure than the second.

The dangerous next second

Picture a large dog asleep beside its owner. The leash is intact. The dog's breathing is slow. Current sensory evidence is reassuring.

An anxious system can accept every fact and remain mobilized, because the feared object is not the dog as it is. It is the dog one second from now. Sleeping dog becomes lunging dog. Intact leash becomes broken leash. A calm present contains the seed of a dangerous transition.

This explains why “look around; nothing is happening” can be accurate and ineffective. The system may not be disputing the present. Its objection is that the present has not guaranteed the next moment.

The same logic can migrate indefinitely. A normal medical result answers a question about what was measured under particular conditions. Anxiety asks what might develop later, what the test might have missed, or whether a different test would close the remaining space. A partner's reassurance describes a current intention. Anxiety asks whether intentions can change. A well-maintained aircraft is safe by ordinary standards. Anxiety asks whether a rare failure remains possible.

In each case, evidence about now is asked to underwrite every version of next. It cannot.

The smoke detector, used with care

Threat detection operates under unequal costs. Missing a real fire may be catastrophic; responding to burnt toast is usually only inconvenient. Under such conditions, a system biased toward some

false alarms can be adaptive. The point is not that evolution designed an ideal smoke detector in the skull. The point is that the threshold for defensive action depends on what false positives and false negatives cost.

The analogy becomes misleading when stretched too far. A smoke detector has one narrow variable to monitor. Human beings anticipate social loss, pain, guilt, illness, uncertainty, abandonment, violence, financial collapse, and threats to people they love. We learn through language and culture. We inherit social conditions. We can fear being afraid. We can imagine an event, judge ourselves for imagining it, and then fear that the judgment reveals something dangerous about our character. No household appliance has this problem.

Nor does evolutionary usefulness imply present-day optimality. A bias can be understandable without being helpful in its current form. It can also be entirely appropriate under real danger. Any account that calls vigilance “disordered” without examining the environment risks locating the problem inside the person when the environment is doing the threatening.

The disciplined version of the smoke-detector principle is modest: defensive thresholds reflect the expected costs of different errors. To understand anxiety, we must ask not only how likely a person thinks harm is, but what they believe it would mean to miss it.

Possibility is not probability

“But it could happen” is one of anxiety's strongest arguments, partly because it is so often true.

The plane could crash. A relationship could end. A presentation could go badly. A symptom could indicate illness. A parent could fail to prevent a child's pain. A morally careful person could make a consequential mistake.

Mere possibility is a poor place to conduct the debate. Almost any open future contains adverse branches. The useful questions begin after possibility has been conceded:

How probable is this outcome under these conditions?

How costly would it be?

What can realistically be controlled?

What evidence would change the estimate?

How capable might the person be of responding if it occurred?

And what is the cost of organizing life around preventing it?

These questions do not guarantee calm. Sometimes a low-probability event deserves serious preparation because its cost is immense. Sometimes a common risk is voluntarily accepted because avoiding it would make valued life impossible. Risk cannot be reduced to a single number. But probability, cost, controllability, recoverability, and opportunity cost belong in the same conversation.

Anxiety can distort that conversation by allowing possibility to carry the entire burden. Once “could” is enough to compel protection, the future contains an unlimited supply of commands.

Unequal standards of proof

A striking asymmetry appears in many forms of severe anxiety:

Threat needs only to remain possible before defence feels justified.

Safety must approach certainty before movement feels permitted.

Consider a flight. A vivid scenario may be enough to cancel even when the person does not believe a crash is likely. To board, however, they may feel they need proof that a crash cannot occur. The evidentiary standards are not simply cautious; they are structurally unequal.

In health anxiety, one anomalous sensation may reopen a catastrophic hypothesis, while repeated normal evaluations are treated as temporary statements about the past. In a relationship, one delayed reply can count as fresh evidence of abandonment, while years of care are discounted because they cannot guarantee

tomorrow. In obsessive doubt, the smallest residue of uncertainty may outweigh a substantial history of responsible action.

This does not mean that tests, reassurance, or careful checking are always inappropriate. Symptoms sometimes require evaluation. Relationships sometimes change. Safety procedures matter. The issue is whether evidence is allowed to update the threatening model, or whether the model has adopted a rule under which disconfirmation is impossible.

A theory that can explain every outcome predicts nothing. A personal threat model can acquire the same defect. If bad news confirms danger and good news confirms only that danger has not appeared yet, the model cannot lose.

The branching future

Imagine next month as a branching tree. One branch contains a new job. From it extend further branches: interest, boredom, friendship, error, competence, conflict, growth, embarrassment. Another branch contains staying in the current role, with a different set of gains and losses. Real decisions are forests, not coin tosses.

Anxiety can change the weighting of the tree without erasing its positive branches. The new job may remain deeply desirable. Yet one branch-public failure-can acquire enough subjective cost to veto the whole node from which it grows.

A journey is declined not because every outcome is bad, but because one branch contains panic. Intimacy is avoided not because connection has no value, but because one branch contains abandonment. A medical examination is postponed not because health is unimportant, but because one branch contains intolerable news.

This is different from a future in which nothing seems worth pursuing. Positive possibilities can retain their value and still become inaccessible because threatening possibilities have excessive stopping power.

The Flow Hijacked synthesis will later describe this as a possible compression of reachable futures. That phrase is not a finding from a scanner or a diagnostic criterion. It is a proposed way to measure what symptom counts often miss: not only how bad a person feels, but how many valued paths remain enterable under uncertainty.

Real danger and calibrated anxiety

Any serious account of anxiety must leave room for the world to be dangerous.

Some people live with war, violence, unstable housing, discrimination, precarious employment, serious illness, or responsibility for vulnerable others. Hypervigilance can be shaped by conditions in which signals have often mattered. A model that treats all threat expectation as internal error would be scientifically crude and morally careless.

Calibration is relational. It concerns a person in an environment, with a history, body, social position, and available resources. The same vigilance may be excessive in one setting, protective in another, and difficult to recalibrate after conditions change. A threat system does not receive an official notice when the statistical structure of life has shifted.

The aim is not fearlessness. A good threat system should interrupt when interruption is warranted, learn when conditions change, and release control when other actions become more valuable. It should be sensitive without being sovereign.

From forecast to action

So far, we have treated anxiety as a way possible danger enters the present. But a forecast matters biologically only if it changes what the organism does.

The next step is therefore behavioural. A prediction of humiliation reduces speech. A prediction of panic selects the seat near the exit. A prediction of illness directs attention inward. Each action may be

sensible. Each action also changes the evidence that becomes available afterward.

At that point, anxiety is no longer only anticipating the future. It is helping to build the future from which it will learn.

Chapter 3 - When Prediction Becomes Policy

When Prediction Becomes Protection

A forecast is an action proposal

A weather forecast that never influences clothing or travel is decorative meteorology. Pain that never alters movement has lost much of its protective function. Prediction matters to living systems because it changes policy: approach, wait, hide, inspect, ask, prepare, escape.

Anxiety becomes durable at precisely this crossing from forecast to action.

Predict embarrassment and speak less. Predict panic and sit near the door. Predict illness and inspect the body. Predict regret and postpone the decision. Predict abandonment and ask for reassurance. These actions may reduce danger, distress, or both. They also alter the next observation.

That final point changes the story. We often picture learning as a clean exchange in which the world produces data and the mind updates its beliefs. Real organisms are active samplers. They turn toward some signals and away from others. They leave rooms, search websites, ask questions, avoid elevators, carry medicines, reread messages, and select companions. Beliefs do not merely receive evidence. They help choose which evidence will exist for the learner.

This is ordinarily intelligent. A physician orders a test because of a hypothesis. A scientist points an instrument where a signal is expected. A hungry person opens the cupboard. Active sampling is how finite organisms avoid examining everything.

It also creates a vulnerability. A model can select evidence in ways that make the model difficult to revise.

The party that remains imaginary

Suppose someone expects humiliation at a party and stays home.

The decision produces genuine information. Staying home reduces immediate uncertainty. Distress may fall. The person learns that cancellation is available and that it changes internal state. What they do not learn is what would have happened at the party.

The feared counterfactual remains untouched. It can still contain awkward silence, visible shaking, rejection, panic, or a spectacular social error. Reality has not shown those outcomes to be false. Nor has it shown them to be true. The branch was never entered.

Repeat this pattern across many situations and the dataset becomes uneven. There may be hundreds of observations of life under protection and very few observations of life without it. The person accumulates confidence that avoidance brings relief while the danger of approach remains estimated from memory, imagination, and rare adverse episodes.

This is not simply “failing to face a fear.” It is a problem of missing counterfactual evidence. The action changes what can be learned.

Relief is a precise teacher

Protection often produces an immediate consequence:

Cancel the appointment, and tension falls.

Leave the crowded room, and breathing eases.

Check the lock, and doubt quiets.

Ask “Are we okay?”, and the relationship feels stable for a moment.

The relief is not fake. It is one of the clearest pieces of feedback in the entire loop. An action has changed the organism's state, sometimes within seconds. Behaviour that reliably terminates an aversive state becomes more likely when a similar state returns. This is negative reinforcement in its technical sense: the removal or reduction of something unpleasant strengthens the action that preceded it. “Negative” describes subtraction, not moral value.

The speed of the lesson matters. Long-term costs are delayed, diffuse, and difficult to attribute. Immediate relief is vivid. The nervous system does not need a lecture on reinforcement to notice which button makes distress stop.

That is why intellectual knowledge and action policy can diverge. A person can know that repeated reassurance maintains doubt and still ask again. They can understand flight statistics and still cancel. They can recognize that checking consumes hours and still feel, at the moment of threat, that not checking would be reckless. The verbal model and the action history are related, but they are not identical.

Who gets credit for safety?

Now suppose the person enters the feared situation while carrying a dense architecture of protection. They fly only after checking turbulence forecasts repeatedly, choosing an aisle seat, holding a companion's hand, taking prescribed medication, monitoring every sound, and keeping an escape plan in mind. The aircraft lands normally.

What caused the safe outcome?

Outcomes do not arrive with causal labels. One interpretation is that the catastrophic estimate was too strong. Another is that the precautions prevented catastrophe or made survival possible. If the second interpretation dominates, a successful flight can strengthen the perceived necessity of the protective system.

Experimental work illustrates this credit-assignment problem. In a human conditioning study, an avoidance response that ostensibly protected participants from threat helped preserve threat beliefs under laboratory conditions. ^[LE-06] The result supports a bounded claim: some safety behaviour can interfere with updating when it is understood as the reason harm did not occur. It does not establish that every precaution is harmful, that all comfort must be removed, or that protection has the same function for every person.

Medication taken as prescribed is not a moral failure. Choosing an aisle seat is not a diagnosis. Holding a hand can enable valued action.

The scientifically useful question is functional: does a support make contact with new evidence possible, or does it make the feared model immune to that evidence? The same object can do either in different contexts.

Local intelligence, global restriction

Avoidance is often called irrational because its long-term costs are easy to see from outside. From inside the moment, it may be the most reliable regulation strategy available.

Take the stairs and fear falls. Cancel the difficult conversation and the immediate threat disappears. Search for reassurance and uncertainty contracts. Each policy has local intelligence. It was selected because it worked at something.

The trouble emerges across timescales. The person loses opportunities to learn that distress can change without escape, that feared outcomes are less probable than expected, that coping is possible, or that mistakes are survivable. Skills remain unpractised. Safe territory contracts. The protective policy becomes fluent while alternatives remain awkward.

Laboratory and clinical research on costly avoidance offers a limited but instructive window. In one relatively small study comparing people with anxiety disorders and controls, some patients continued acquired avoidance even when it sacrificed competing rewards; importantly, the patient distribution was not uniform. ^[LE-07] That heterogeneity matters. It prevents a story in which every anxious person has one fixed avoidance defect. Work using latent learning trajectories likewise finds multiple patterns within groups rather than a single conditioning signature. ^[LE-08]

Empirical result: avoidance can persist despite cost in some people and contexts, but there is substantial heterogeneity. **Interpretation:** a policy optimized for rapid relief may degrade longer-term adaptability. The second statement is a systems-level reading of the evidence, not something one experiment directly measured.

Environmental engineering

Protection can expand beyond individual actions into the design of a life.

Always sit near exits. Travel only with one trusted person. Choose routes that pass hospitals. Avoid heat, caffeine, crowds, conflict, unfamiliar food, exercise, spontaneity, public speaking, or being far from home. Arrange work so evaluation is rare. Ask family members not to mention certain topics. Keep every day predictable enough that surprise has little room to enter.

Any one of these choices may be reasonable. Some may be medically indicated. Some may simply reflect preference. The question is cumulative and functional: what happens to the person's learning environment when protection becomes architecture?

A tightly controlled environment can reduce symptoms because triggers are absent. That improvement may be welcome and necessary. But symptom reduction produced by compression is not identical to increased flexibility. A person with fewer panic attacks because they never leave home has reached a different state from a person with fewer panic attacks while travelling, exercising, working, and entering uncertainty.

Outcome measures that count episodes without measuring reachability can confuse the two.

Protection from information

We can now name a broader interpretation. Anxiety may protect a person not only from feared outcomes but, unintentionally, from information that could revise the forecast.

Avoidance prevents sampling. Rituals can claim credit for safety. Reassurance replaces uncertain observation with another person's temporary conclusion. Hypervigilance selects anomalies while ordinary non-events pass unrecorded. Environmental control reduces variance so thoroughly that flexibility is never tested.

None of this proves that the world is safe. Corrective evidence is not synonymous with pleasant evidence. Entering a situation may reveal a real problem. A conversation may go badly. A symptom may warrant treatment. What matters is whether the system can gather sufficiently representative information to calibrate its model, rather than sampling almost exclusively through the policy it is trying to evaluate.

This is a difficult standard because no person samples the world neutrally. All perception and action are selective. The goal is not omniscience. It is enough openness that a threat prediction can, in principle, be weakened by what happens.

The first Flow Hijacked proposition

The Flow Hijacked synthesis proposes the following pattern: under anticipated threat, protective policies may narrow action; narrowed action may reduce access to disconfirming and coping-relevant evidence; the resulting evidence environment may preserve the forecast that protection is necessary.

This is a synthesis, not an established mechanism. Its nearest evidence comes from work on avoidance, safety behaviour, generalization, and reinforcement, but no study has validated “Flow Hijacked” as a clinical construct. ^{[LE-06][LE-07]}

The proposal becomes scientifically useful only when it risks being wrong. If avoidance does not predict later contraction of valued action beyond symptom severity and environmental constraint, the claim weakens. If people can substantially expand reachability without changing the relevant protective policies, the causal story needs revision. If reduced opportunity reliably precedes avoidance rather than follows it, the direction may be different from the one proposed.

For now, the model remains backstage. Older learning science must first explain how a previously neutral cue becomes a warning, why a safe outcome may not erase that warning, and how context decides which history governs the next action.

Prediction has become protection. The next question is how protection learned what to protect against.

PART II

LEARNING DANGER

What threat, safety, memory, and context teach the nervous system.

Chapter 4 - How Fear Learns

How Fear Learns

When a sound acquires a future

A tone sounds. Nothing follows.

Later, the tone sounds and an unpleasant shock arrives. After enough pairings, the tone alone changes the body. Skin conductance rises. Attention sharpens. The cue has acquired a future.

The laboratory arrangement is austere on purpose. Human lives do not arrive as coloured shapes followed by carefully calibrated electrical stimulation. By stripping a learning episode down, however, the experiment reveals a basic capacity: an organism can learn that one event predicts another and begin preparing before the second event occurs.

Preparation at the cue is useful. Waiting for harm would squander the warning. Fear learning turns sequence into anticipation.

The familiar word “association” can make this sound like two items being tied together in storage. Prediction is more revealing. The cue matters because of what it says about what comes next. A hospital corridor may predict news. A notification tone may predict conflict. A facial expression may predict rejection. A flutter in the chest may predict a cascade of sensations. The physical cue has not necessarily changed. Its temporal meaning has.

Human beings acquire such meanings by more than direct injury. We learn by observing other people, receiving instructions, absorbing cultural narratives, remembering prior events, and imagining consequences. A child may avoid dogs after a bite, after seeing a caregiver recoil, or after hearing repeated warnings. These routes are not interchangeable, and none guarantees a lasting fear. They show that a nervous system embedded in language and relationships can inherit forecasts without personally encountering every outcome.

Prediction is not a photograph

A learned threat is not stored as a perfect image of one event. Learning must compress.

If a bite taught only that one particular dog, at one angle, on one street, at one time of day was dangerous, the lesson would be too narrow to protect. The learner extracts features and relations: large dog, sudden movement, exposed teeth, lack of control. A humiliating talk may teach something broader than “that audience disliked that sentence.” It may teach that speaking under evaluation is dangerous.

This compression is adaptive because the world does not repeat itself exactly. Tomorrow's danger will differ from yesterday's. A system must decide which differences are irrelevant and which make a new situation safe.

That decision has no perfect solution. Generalize too narrowly and related danger is missed. Generalize too broadly and safe situations are recruited into defence. Later we will examine what happens when one danger becomes many. Here the important point is earlier: fear learning produces a rule, and rules require a boundary.

Two lessons at once

Pavlovian learning concerns predictive relations among events: this cue has been followed by that outcome. Instrumental learning concerns what actions produce: when I leave, distress falls; when I check, uncertainty briefly contracts; when I remain silent, evaluation becomes less likely.

In anxiety, these lessons can couple.

The cue says, “Danger may occur here.”

The action says, “Escape changes the state.”

Relief says, “Use that action again.”

The system has now learned both a forecast and a response. Catastrophe need not occur repeatedly for the pattern to persist. The

cue can continue to evoke threat, while the relief following protection continues to rehearse the policy.

This is why an old fear can survive long periods without the feared event. Absence is ambiguous when the person repeatedly performs an action believed to prevent harm. The world did not deliver catastrophe, but the learning system still has to decide why.

What conditioning studies actually show

The scientific literature does not support one simple claim that “anxious people condition more.” An updated meta-analysis of classical fear-conditioning studies found group differences particularly in responding to safety cues and during extinction, but the 44 included studies varied in diagnoses, tasks, and measures. ^[LE-01] Earlier meta-analytic work also identified group-level conditioning differences, but it drew on an older and more limited literature. ^[LE-02]

The heterogeneity is not a footnote. Analyses of learning trajectories have found distinct patterns within diagnostic groups, with some potentially dysfunctional trajectories overrepresented among patients rather than universally present. ^[LE-08] A person can experience severe clinical anxiety without displaying a dramatic abnormality in one laboratory task; another may show a particular learning pattern that does not summarize the rest of their life.

Empirical result: on average, anxiety-related groups differ from controls on some aspects of threat and safety learning, especially in certain phases and measures, but no single conditioning profile characterizes everyone. ^{[LE-01][LE-08]}

Interpretation: clinical anxiety may arise through multiple learning routes and may be maintained by different combinations of acquisition, discrimination, generalization, avoidance, and retrieval. That plural account fits the evidence better than a universal broken-fear-learning mechanism.

The laboratory's value lies in controlled questions, not miniature biographies. A skin-conductance response to a cue is not a year of social anxiety. A startle measure is not the meaning of illness, shame,

or responsibility. Experimental tasks isolate components. Clinical lives reveal how components combine with language, bodily state, development, culture, and real consequence.

There is no lone fear switch

Fear learning is often illustrated by an image of the amygdala glowing red, as though the rest of the organism were waiting for instructions. That picture is memorable and incomplete.

A meta-analysis of human fear-conditioning imaging studies found coordinated involvement across systems related to salience, interoception, autonomic regulation, and action, alongside decreases in other regions; it did not discover a solitary fear centre. ^[NB-03] The amygdala matters in many threat-learning contexts, but “amygdala equals fear” collapses a distributed process into a mascot.

Brain imaging also measures group-level changes in blood-oxygen signals, not private fear itself. An activated region may participate in many tasks. A difference between conditions does not prove that the region caused the subjective state, and it does not provide an individual diagnostic signature.

Broader scientific inference: threat learning is implemented by interacting brain, bodily, and behavioural systems, with the relevant configuration changing by cue, task, context, and person. This inference is more defensible than the story of one internal alarm issuing uniform commands.

The coordinates hidden in every lesson

Suppose a crowded room is tolerable with a trusted friend but frightening alone. A rapid heartbeat is welcome during exercise and alarming while sitting at a desk. A dog feels safe beside its attentive owner and dangerous when loose. The cue does not carry one context-free meaning.

Learning includes hidden coordinates: location, bodily state, social company, time, recent events, and perceived control. The system may not learn “elevators are safe.” It may learn something closer to “this

elevator, with this person, while rested, near an exit, produced this outcome.”

Hidden conditionality becomes visible when learning fails to travel. Someone manages a feared situation in a therapist's office or with a companion, then feels astonished when anxiety returns elsewhere. The new learning was real. It may also have been narrow.

This is not a reason to dismiss supported practice as artificial. Scaffolding can make learning possible. It is a reason to ask what, exactly, was learned and which features the learner treated as essential.

Flow Hijacked at the level of learning

The Flow Hijacked synthesis does not propose a new kind of conditioning. It asks what happens when ordinary predictive and instrumental learning repeatedly favours one narrow family of protective actions.

One possibility is that threat cues come to activate not only a feared outcome but a small, well-rehearsed policy set: inspect, escape, postpone, obtain certainty. Because those actions change evidence and internal state, the system gains repeated experience with the policy and limited experience beyond it. Over time, the action landscape may become easier to enter in one direction and harder in others.

That is a hypothesis about interaction and accumulation. It is not demonstrated by conditioning meta-analyses, and it should not be retrofitted onto every anxious life. A different person may be constrained mainly by external danger, depression, illness, exhaustion, lack of opportunity, or another process entirely.

What learning science securely gives us is both simpler and more consequential: cues can acquire predictive threat value; actions can become learned solutions; learning generalizes; and context helps decide which lesson is retrieved.

One question remains before the story can move forward. If danger is learned, what does the organism have to learn when danger does not arrive?

Chapter 5 - Safety Is Learned Too

Safety Is Learned Too

The quiet half of survival

Fear attracts attention because fear is loud. Safety learning is quieter and equally necessary.

An organism that detects danger but never learns when danger has ended cannot allocate itself well. It needs to discover when a warning no longer applies, when a situation can be entered, when defensive mobilization can be released, and which conditions distinguish danger from safety.

It is easy to imagine safety as the default left behind when nothing bad happens. Learning is less obliging. If a cue has predicted harm, one harmless encounter may not remove the old relation. The system must infer what the absence means.

Perhaps the cue is now safe.

Perhaps this occasion was an exception.

Perhaps protection worked.

Perhaps danger remains but was delayed.

Safe outcomes do not interpret themselves.

Extinction is not a delete key

In conditioning research, extinction usually means repeatedly presenting a previously threatening cue without the expected aversive outcome. Defensive responding often declines. The word unfortunately invites the image of the old learning being extinguished like a flame.

Evidence argues for more caution. Responding can return after time passes, after a context changes, or after an aversive event occurs again. Foundational animal studies demonstrated reinstatement and contextual renewal, showing that responding after extinction remains sensitive to later events and retrieval conditions. ^{[LE-09][LE-10]} These

experiments establish principles in controlled animal learning; they do not prove that a person's clinical recurrence is caused by the same mechanism.

Human imaging research on extinction also implicates distributed, condition-dependent systems, including prefrontal, cingulate, insular, hippocampal, and amygdala-related processes. ^{[LE-05][NB-04]} The evidence does not support a cartoon in which one cortical “brake” simply switches off one subcortical alarm.

Empirical result: extinction learning can reduce conditioned responding, while old responding can return under certain conditions; human extinction engages distributed systems. ^{[LE-09][LE-10][NB-04]}

Interpretation: extinction often creates new learning that competes with, rather than literally erases, prior learning. The phrase is a strong organizing interpretation, not a claim that every memory contains two neatly separable files.

Improvement can therefore be genuine without being invulnerable. Old learning surviving is not the same as new learning failing.

The elevator experiment

Imagine a person who predicts that intense anxiety in an elevator will cause collapse or loss of control. They enter with a clinician. The doors close. Anxiety rises to eight out of ten and stays high. The elevator reaches the chosen floor. The person steps out shaking.

Was the exercise unsuccessful?

Not necessarily. If the critical prediction was “this degree of arousal will make me collapse,” the observed outcome may be informative even without a graceful decline in distress. The person remained upright, completed the journey, and discovered something about the relation between feeling and outcome.

Now imagine a second attempt. Anxiety falls rapidly because the person checks pulse every few seconds, grips emergency medication, stands beside the controls, and receives continual reassurance. The

graph of distress looks excellent. Yet the learner may conclude, “I was safe because I prevented disaster.”

Was this attempt unsuccessful?

Again, not necessarily. The supports may have made entry possible and may form a sensible first step. The point is that symptom reduction alone does not tell us what was learned.

The central question is causal: why did the feared outcome not occur?

Expectancy violation without theatrical suffering

Modern exposure accounts have emphasized expectancy violation and inhibitory learning. What did the person predict? What happened? Was the discrepancy noticed? Did the experience change what is expected next time?

If the prediction is “I cannot function until anxiety falls,” completing a valued action while anxiety remains may be more informative than rapid habituation. If the prediction is “visible uncertainty will lead to rejection,” an imperfect conversation followed by continued connection can test it. If the prediction is “doubt must be eliminated before I decide,” making a reversible decision with some doubt intact creates relevant evidence.

None of this makes maximal distress the goal. Exposure is sometimes caricatured as forcing a person to endure the worst until fear surrenders. That is neither a humane principle nor a scientifically necessary one. Learning requires engagement with the relevant prediction, not a performance of suffering.

Overwhelming intensity can prompt escape, dissociation, or a conclusion that the event was indeed intolerable. Ethical work preserves informed choice, pacing, genuine safety boundaries, and collaboration. An exercise should be designed to learn, not to prove obedience or courage.

This distinction matters outside formal therapy as well. People often conduct accidental exposure exercises in daily life. A difficult meeting,

a train ride, a night alone, or a bodily sensation becomes an experiment whether or not anyone names it. The resulting lesson depends on prediction, attention, action, attribution, and context.

Safety behaviour has no moral category

Some safety behaviours can preserve threat beliefs in experimental conditions when the behaviour is credited with preventing harm. [LE-06] It does not follow that all support is contamination.

A handrail may help someone enter a staircase they otherwise cannot approach. A trusted companion may make a first journey possible. Carrying appropriate medication may be medically sound. The relevant question is not whether an object appears on a prohibited list; it is what function the behaviour serves in this learning episode.

Does it expand contact with the feared uncertainty, or does it prevent the key prediction from being tested?

Can it later become more flexible, or has it become a non-negotiable condition of movement?

Does the person attribute safety to a more calibrated understanding of the situation and their ability to respond, or exclusively to the ritual?

The same support may be scaffold in one phase and cage in another. Function changes over time.

Safety can mean more than “nothing bad happens”

No honest treatment can teach that nothing bad ever occurs. Flights are not metaphysically guaranteed. Bodies become ill. People sometimes judge one another. Relationships can end. A doctrine of universal safety would collapse at first contact with life.

Corrective learning can concern several different quantities. A feared outcome may be less likely than predicted. Its cost may be lower. Coping may be greater. Recovery may be faster. Uncertainty itself may be tolerable. A person may discover that anxiety is unpleasant without being an instruction.

This widens the meaning of safety. It is not only the non-occurrence of harm. It can be the discovery that one remains capable of action, adaptation, repair, and return.

Such learning is often visible less in what a person reports during an exercise than in what becomes possible later. They enter again. They stay without continuous monitoring. They vary the route. They recover after a difficult experience. A feared feeling may still arrive, but it no longer possesses exclusive authority over behaviour.

Safety must travel

A person may learn that one flight was tolerable and still fear a different airline. One social event may go better than expected while a professional gathering remains inaccessible. Palpitations during supervised exercise may feel safe, while the same sensation after a sleepless night feels ominous.

The nervous system asks, implicitly, whether the present belongs to the same category as the learning episode. If the new lesson was encoded narrowly, it may govern only a small region.

Variation can help clarify what is essential and what is incidental: different contexts, different bodily states, different times, different levels of support. But variation must remain proportionate and safe. No person can rehearse every possible future, and no responsible account should invite experimentation with medically concerning symptoms without evaluation.

The goal is not exhaustive exposure to life. It is sufficiently abstract learning: I can approach uncertainty; I can notice prediction; I can remain in contact with evidence; I can respond if difficulty occurs.

What the Flow Hijacked synthesis adds

The Flow Hijacked interpretation treats safety learning not merely as a decline in fear but as a potential reopening of action. A safe or survivable experience matters when it changes which paths remain reachable later.

This is a broader outcome than immediate distress. It can also fail independently of distress. A person may feel calmer while life remains compressed, or feel anxious while their field of action expands.

That proposal is not an established measure. It creates a research obligation: define reachability, observe it over time, separate it from mood and resources, and test whether particular learning experiences predict durable expansion. If they do not, the synthesis must change.

For the moment, the evidence supports a more modest lesson. Safety is learned in relation to cues, contexts, actions, bodies, and outcomes. It is not bestowed by a slogan. And because every lesson carries context, new learning faces a second challenge after it is acquired.

It has to be found again.

Chapter 6 - Memory, Context, and the Return of Fear

Memory, Context, and the Return of Fear

The swimming lesson problem

A person learns to swim in a warm indoor pool. The water is clear. An instructor stands nearby. The edge is always within reach.

Months later, the same person enters a cold sea. The bottom is invisible. Waves interrupt breathing. The shore looks farther away from the water than it did from the sand.

“But you know how to swim” is technically true and psychologically incomplete.

The skill was acquired. Whether it governs action here depends on retrieval, context, bodily state, and the degree to which learning travelled beyond its original coordinates. The sea has not removed the lesson. It has changed the conditions under which the lesson must be found.

Anxiety treatment and ordinary recovery repeatedly encounter this problem. Someone understands panic in a calm office and loses access to that understanding on a crowded platform. A conversation feels manageable after rehearsal and impossible when a particular person enters the room. A period of progress survives ordinary days and falters after three nights of poor sleep.

The knowledge did not necessarily disappear. Deployment changed.

Context is not the wallpaper

We commonly use “context” to mean place: office, home, airport, supermarket. Place matters, but the nervous system's context is larger.

The body is context. Fatigue, pain, illness, hunger, intoxication, withdrawal, hormonal change, dehydration, and stimulation can alter the state in which a cue is encountered. The heart racing during a

chosen run does not have the same meaning as the heart racing in bed. The sensory event overlaps; the surrounding model differs.

Social company is context. Alone is not accompanied. With a trusted person is not among strangers. Being observed is not being anonymous. A relationship can alter both actual safety and the expected ability to cope.

Time is context. Night retrieves different concerns from midday. The anniversary of an event may change salience. Long gaps can weaken access to a recently learned response. Recent stress changes what feels near.

Control is context. A door that could be opened feels different from an identical door believed to be locked. Choosing to enter differs from being compelled. The action may look the same to an observer while the learner inhabits a different contingency.

Context is therefore not decorative wallpaper around the “real” cue. It is part of the lesson's address.

Three ways an old response can return

Learning research uses several terms for the return of responding after extinction.

Renewal refers to responding that returns when context changes after extinction.

Reinstatement refers to responding restored after renewed exposure to an aversive event.

Spontaneous recovery refers to responding that reappears with the passage of time.

Classic animal experiments helped establish contextual renewal and reinstatement. ^{[LE-09][LE-10]} These findings matter because they contradict a simple deletion model. They do not supply a one-to-one explanation for clinical recurrence. A laboratory cue, an aversive stimulus, and a controlled context are far simpler than a human life shaped by language, relationships, illness, values, and years of learning.

Human extinction imaging is consistent with the involvement of multiple systems in learning and recalling extinction, including context-sensitive hippocampal and prefrontal processes, but these studies are largely conducted in healthy participants and remain correlational at the level of blood-oxygen signals. ^[NB-04]

Empirical result: conditioned responding can return after extinction, and retrieval depends partly on time, later aversive events, and context in controlled paradigms. ^{[LE-09][LE-10]}

Broader scientific interpretation: improvement often involves new learning whose influence varies with retrieval conditions rather than an irreversible erasure of prior learning.

Boundary: a recurrence of anxiety in a person cannot be diagnosed as “renewal” merely because it followed a change of setting. Similar experiences can arise from new danger, medical changes, sleep loss, stress, treatment interruption, or other causes. The laboratory vocabulary offers questions, not verdicts.

When the body resembles an old place

Consider someone who has learned, through repeated and appropriately evaluated experiences, that palpitations can occur without catastrophe. Most days, the new understanding holds.

Then comes a night after travel: little sleep, too much coffee, dehydration, unfamiliar surroundings, no trusted person nearby. The heart accelerates. The current bodily state resembles earlier panic episodes more than it resembles the calm conditions under which much of the newer learning occurred.

An older interpretation may become easier to retrieve.

It is tempting to tell the person, “You already know this.” That statement can add shame to fear. Knowing is not one switch. A proposition available during calm reflection may lose competition when bodily and contextual cues change. Retrieval is an achievement performed in a state, not a permanent possession displayed equally everywhere.

This is also why a body should not be treated merely as noise around cognition. Bodily change can be medically meaningful. New, severe, or concerning symptoms deserve appropriate assessment rather than automatic classification as anxiety. Once relevant safety questions have been addressed, bodily context still matters to learning because it changes which model seems plausible and which action feels available.

Memory does not return the past untouched

The past does not sit in storage as a complete recording. Remembering is constructive. Details are selected and reassembled in light of current goals, meanings, cues, and knowledge.

This does not mean memories are arbitrary or that painful events did not happen. It means recollection is an active biological process rather than the playing of an internal video. A remembered panic episode may become organized around “the time I nearly collapsed,” even if collapse did not occur. A social encounter may be recalled as “the night everyone noticed,” while the actual evidence was mixed. The emotional conclusion becomes the title under which the event is filed.

Later, the titled memory serves as evidence for a forecast. Past and future collaborate.

Memory supplies candidate examples. Imagination recombines them. Current bodily state lends some examples greater plausibility. Attention selects confirming details. Action determines which new memories will be created.

This recursive picture is a broader scientific interpretation, not a claim that memory is always distorted toward threat or that recollection alone causes anxiety. Memory can correct fear as readily as preserve it. A person may remember coping, repair, help received, or a catastrophe that did not occur. The issue is which episodes become accessible and how they are used.

Why recurrence feels like betrayal

Few experiences are more demoralizing than the return of an old state after genuine progress.

“I thought I was past this.”

The sentence can quickly become an accusation. The treatment failed. The progress was fake. I have returned to the beginning. Something about me is permanently broken.

A learning account permits a more accurate and less punitive reading. New learning may have become influential across many conditions and then lost retrieval strength in a particular state. The old response returning does not erase the days, choices, and capacities that emerged after it. Nor does prior improvement guarantee that the present recurrence is trivial.

Both truths can stand. The person has changed, and the old pattern is active now.

This changes the next question. Instead of asking only why the fear came back, one can ask what made this situation a good retrieval context for the older policy. Was there a new aversive event? A shift in sleep, health, medication, substance use, social safety, perceived control, or cumulative stress? Had practice become confined to one setting? Was the feared outcome never fully tested because a ritual retained causal credit?

The questions turn recurrence from a verdict into information. They do not promise an easy answer, but they restore curiosity where shame would otherwise close inquiry.

Acquisition is only the first problem

Change faces at least two learning problems.

The first is acquisition: can something new be learned?

The second is deployment: will the learning influence behaviour when the world, body, social field, or time differs from the learning episode?

Deployment is routinely underestimated because it hides during successful sessions. A person can articulate a nuanced understanding in daylight and struggle to access it at three in the morning. They can complete an exercise with a clinician and avoid the same situation alone. They can tolerate uncertainty in one domain while remaining rigid in another.

None of this means the original work was useless. It shows that portability is part of the work.

Portable learning is not simply the repetition of reassuring sentences in more places. It may require encounters with meaningful variation, opportunities to attribute outcomes differently, practice recovering after difficulty, and a less conditional sense of agency. The abstraction that travels is often not “this exact situation is safe.” It is “I can remain in contact with evidence, act under uncertainty, and respond when conditions change.”

The therapist, companion, and borrowed context

Support can change what is possible in the moment. Experimental studies suggest that familiar support figures can attenuate acute threat-related distress or responding under some conditions, though effects depend on relationship, age, task, and sample. ^{[SR-03][SR-04][SR-07]}

^[SR-08] The evidence concerns bounded experimental outcomes, not proof that another person's nervous system literally becomes a portable safety device.

Support creates a valuable paradox. It can open exploration, allowing a person to enter a situation and encounter evidence that would otherwise remain unavailable. It can also become part of a very narrow context: “I can do this only when this person is here.”

The distinction is not solved by withdrawing support on principle. It is solved by attending to function and transfer. Does the relationship enlarge action and gradually support wider agency, or does every safe outcome get credited solely to the companion's presence?

The Flow Hijacked synthesis sometimes uses the phrase “borrowed safety” for the first possibility: connection temporarily widens the field

enough for exploration. This is a narrative and research proposal, not a validated neural mechanism. If support lowers distress without increasing approach, exploration, or later independent reachability, the stronger version of the proposal would fail. The studies cited here support acute social regulation, not the full Flow Hijacked claim.

A map that changes with the traveller

The world is not divided once and for all into safe and dangerous locations. Its reachable map changes with the state of the traveller.

A route available while rested may disappear under exhaustion. A conversation possible with support may close when alone. A bodily sensation interpreted with curiosity in one setting may become an emergency signal in another. The external geography remains; the action geography changes.

This is where the Flow Hijacked lens becomes useful but must remain disciplined. **Flow Hijacked synthesis:** anxious recurrence may sometimes reflect a context-driven contraction of reachable action, in which older protective policies become more accessible than recently learned alternatives. This is not a finding established by renewal studies. It is a testable integration that would require dense within-person measurement of states, contexts, policies, and subsequent action.

Evidence that would weaken it is easy to imagine. If context changes predict distress but not changes in reachable action, the proposed link is too strong. If external constraints fully explain the contraction, an internal policy account adds little. If the newer policy remains accessible yet is rejected for good reasons, “retrieval failure” is the wrong explanation.

A useful synthesis should make these distinctions sharper, not absorb every outcome.

The bridge to what comes next

We can now see why threat learning is not contained in one memory trace.

A cue acquires predictive meaning. An action becomes a learned solution. A safe outcome demands interpretation. New learning competes with old learning. Context influences retrieval. Memory reconstructs examples. Bodily state changes plausibility. Support changes what can be approached. Action determines which evidence and memories will be available next.

The system is recursive.

The next complication is unavoidable. A lesson about danger rarely remains attached to one exact cue. It travels across resemblance, meaning, and category. A single dog becomes large dogs, barking, parks, unfamiliar houses. One episode of panic becomes crowded rooms, heat, exercise, distance from help, or the sensation of uncertainty itself.

The future does not shrink all at once. Often, it loses territory by generalization.

PART III

WHEN SAFETY SHRINKS THE WORLD

*How useful protection becomes a narrow architecture
of life.*

Chapter 7: When One Danger Becomes Many

The smoke alarm that learns categories

A child is frightened by a dog that breaks free of its lead. The lesson could remain exquisitely specific: this dog, in this park, moving at this speed, with its owner shouting. Nervous systems rarely file experience with such archival precision. They compress. A large animal running toward me may be dangerous. Barking may precede impact. A lead may fail. Parks may contain dogs.

That compression is not a defect. If danger had to be learned from scratch in every new setting, learning would arrive too late to be useful. A gazelle that recognized only the exact lion it had previously escaped would have a short and educational life. Generalization allows an organism to carry a lesson beyond the incident that taught it.

But every useful compression raises a difficult question: how far should the category extend?

The child may become cautious around large, charging dogs. That is easy to understand. The category may then widen to all dogs, then to parks, then to streets where dogs might appear, then to visiting friends who own one. Nothing magical happens at any step. Similarity, meaning and uncertainty do the work. Yet a single event has now acquired many addresses.

Laboratory studies often represent this process with a generalization gradient. One cue predicts an unpleasant event. Later, people encounter cues that resemble it by different degrees. Defensive responding is usually strongest near the learned cue and falls as similarity decreases. The exact shape of that fall matters. A steep gradient preserves discrimination: this cue is dangerous, a nearby one is not. A shallow gradient spreads defensive value more widely.

Meta-analyses find broader fear generalization on average in groups with anxiety-related disorders, but the effect is neither identical across diagnoses nor consistent across every physiological and self-report

measure. Some people show broad gradients, some do not, and a group average is not a personal signature. ^[LE-03] ^[LE-04] The empirical result is therefore interesting and bounded: overgeneralization appears in the literature, but “anxiety generalizes everything” would be a slogan, not a finding.

Meaning travels farther than appearance

Similarity is not only visual. A red circle may resemble an orange circle, but a hospital corridor can resemble a kitchen table if both are places where one waits for consequential news. A manager's short message can resemble a former partner's silence because both leave the social future unresolved. Two sensations can feel related because both have once preceded panic, even if one is a racing heart and the other is dizziness.

Human beings generalize through concepts, language and stories. We do not merely ask, “Does this look like the old threat?” We also ask, often without words, “Could this mean what that meant?” A humiliating presentation may teach danger about public speaking, but it may also teach something wider: visible uncertainty is unsafe; other minds are unforgiving; mistakes alter belonging. Once the lesson is organized at that level, it can travel into dates, classrooms, interviews and ordinary conversation.

This helps explain why well-meant reassurance can miss. “This is a different room” answers physical resemblance. “These are different people” answers identity. Neither necessarily touches the category that has been learned: situations in which my worth is publicly uncertain.

That is an interpretation of how generalization can become personally coherent, not a claim that every fear hides one secret sentence. Learning can remain sensory, relational, linguistic or mixed. Trauma, developmental history, instruction from others and direct conditioning may all contribute. The useful question is not which single route must have occurred. It is what rule the system now seems to be carrying from one situation to the next.

A category can be trained by what it never meets

Generalization does not end when the original learning ends. The category is revised by later encounters, including encounters prevented by avoidance. If every small dog is approached but every large dog is escaped, the system receives different evidence at different points on the gradient. The safe region gains detail; the threatening region remains governed by simulation.

This produces an important asymmetry. Experience can refine a category only where experience is allowed to occur. The unexplored edge is not empty; it is populated by old learning, stories from other people, alarming images and the body's own predictions. Absence of contact does not preserve a neutral hypothesis. It preserves whatever model was already strong enough to prevent contact.

The point is not that direct exposure always corrects fear. A feared outcome can occur, an encounter can be ambiguous, and an experiment can be too overwhelming to teach useful discrimination. What matters is whether the encounter supplies interpretable evidence at a tolerable cost. New learning requires more than presence. It requires a prediction specific enough to be tested and an outcome that is not automatically credited to protection.

This also explains why courage cannot be measured simply by distance from the feared object. A person may stand close while performing an elaborate internal escape, or remain farther away while making a genuine new choice. The relevant unit is the relation among prediction, action and learning.

Safety has to generalize too

Danger enjoys a practical advantage. One painful event can be enough to justify caution. Safety is harder to establish because “nothing happened this time” leaves room for qualifications. Perhaps the threat was absent only because I was accompanied. Perhaps the dog was calm because its owner held the lead. Perhaps the meeting went well because I rehearsed every sentence. Perhaps the heart settled because I escaped quickly.

New safety learning can therefore remain narrow. A person may manage one lift with one therapist at one time of day and discover, honestly, that this particular lift was tolerable under these conditions. The learning is real. It is simply not yet portable. When the building, companion or bodily state changes, older threat learning may regain control.

Research on conditioning supports the broad idea that extinction learning is often context-sensitive and that extinguished responding can return. Much of the foundational work comes from tightly controlled laboratory and animal paradigms, so it cannot be mapped directly onto a person's clinical course. ^[LE-09] ^[LE-10] The responsible inference is modest: improvement in one context need not mean erasure, and recurrence does not prove that nothing was learned.

This is one reason “But you did it yesterday” can feel so unfair. Yesterday may have supplied different retrieval cues. Sleep, company, location, stress and internal state can alter which history becomes available. The nervous system is not lying when the old response returns. Nor is it revealing an immutable truth. It is retrieving under changed conditions.

The first narrowing

Imagine a map drawn around a threatening point. Generalization enlarges the region that deserves caution. Avoidance then removes some of that region from direct testing. Each unvisited place retains all the futures imagination can give it. The map grows darker precisely where experience becomes scarce.

The Flow Hijacked synthesis begins here, but it must be marked as synthesis rather than established mechanism. We propose that generalization can increase the number of situations in which defensive policies become eligible; avoidance can then reduce the corrective information available in those situations. The result may be a self-maintaining contraction of reachable life. The components are supported in separate literatures. Their assembly into one dynamical architecture is a testable proposal, not a discovered circuit.

This distinction matters because generalization is not the enemy. A life without it would be impossible. The aim is not to make every encounter seem unrelated to every previous encounter. It is to preserve the capacity to discriminate again when the world supplies new evidence.

The next chapter follows the action that most efficiently prevents such evidence from arriving. It is not a foolish action. It is often the action that works first.

Chapter 8: Avoidance, the Solution That Works Too Well

The lift and the staircase

You approach a lift. The doors open. Heat moves through the chest; the air seems suddenly insufficient. The staircase is ten metres away. You turn toward it, and within seconds the body begins to settle.

Something important has happened. A prediction was followed by an action, and the action was followed by relief. The relief is not imaginary. It is a real change in state, often rapid enough to make the sequence memorable. Whatever else the person believes, the nervous system has received a persuasive lesson: when this configuration appears, leaving changes what happens next.

This is negative reinforcement in its plainest form. “Negative” does not mean bad, and “reinforcement” does not mean praise. A behaviour becomes more likely because it removes or reduces something aversive. The staircase has not merely transported a body between floors. It has defeated distress.

Calling the action irrational misses why it returns. Locally, it is excellent. The lift feels dangerous; the stairs reduce fear. A conversation feels dangerous; cancellation removes the immediate social risk. Uncertainty feels unbearable; checking closes it for a moment. The long-term problem is built from short-term successes.

What was not learned

Suppose you stay out of the lift. You learn that avoidance brings relief. You may also learn that you can use it again. What you do not learn is whether the feared event would have happened, whether the sensations would have peaked and fallen, or whether you could have remained present while uncertainty stayed unresolved.

This absence is easily overlooked because non-events do not announce themselves. There is no memory of the catastrophe failing to

occur inside the lift, because there was no journey. The avoided future remains available to imagination in pristine condition.

The same ambiguity can survive partial approach. A person enters the lift but grips a bottle of water, calls a friend, stands beside the controls and exits at the first change in heartbeat. The lift arrives safely. Why? One possible lesson is that the prediction was too severe. Another is that the protective arrangement prevented disaster. Outcomes do not arrive with causal labels attached.

An experimental study showed that an avoidance response could preserve threat beliefs during human fear conditioning. ^[LE-06] That finding does not license the rule that every safety behaviour must be removed. A support can be the bridge that makes new contact with the world possible. The functional question is subtler: does this support enable learning, or does it take exclusive credit for safety?

A handrail is useful if it helps someone enter. It becomes restrictive if the system concludes that entry without it is impossible forever.

When avoidance becomes expensive

Avoidance is easiest to see when someone does not fly, drive or attend a gathering. It also appears in miniature: speaking briefly to prevent scrutiny, sitting near the door, never allowing hunger, scanning every exit, choosing only reversible decisions, arranging conversations so that disagreement cannot occur. A life can remain outwardly busy while becoming internally repetitive.

Research using tasks in which avoidance competes with reward suggests that some people with anxiety disorders continue to avoid even when doing so becomes costly, but the pattern is heterogeneous. In one relevant study the patient distribution was not uniform, and the sample was modest. ^[LE-07] The finding is not that anxious people cannot value reward. It is that, for some people in some states, avoiding threat retains enough value to defeat an available gain.

Over time, the costs can become real rather than merely anticipated. Avoid driving for years and skill may decline. Avoid exertion and fitness may fall. Avoid conversation and social fluency

may weaken. The system then encounters genuine difficulty and treats it as proof that the original prediction was correct. Avoidance has helped manufacture some of the incapacity it feared.

That sentence should never be used to blame. People avoid because the action solves something. Poverty, discrimination, illness, unsafe relationships and actual danger can also narrow action; not every restricted life is an anxiety loop. The task is to identify function and context, not to moralize the visible behaviour.

The architecture of accommodation

Avoidance can recruit other people. A partner makes every phone call. A parent repeatedly confirms that food is safe. A colleague quietly removes presentations. These accommodations can be loving, sensible and temporarily necessary. They can also become part of a system in which one person's relief and another person's care reinforce the same shrinking boundary.

The right question is not “Who is enabling whom?” That language is too blunt for most human relationships. Better questions are: What becomes possible because of the support? What becomes impossible without it? Does support widen exploration over time, or must the protected conditions become increasingly exact?

The distinction between scaffold and cage is temporal. The same arrangement can be one today and the other six months later.

Avoidance is not only staying away

Some avoidance is visible as absence. Much of it occurs while the person is present. One can attend a gathering and never ask a question whose answer matters. One can remain in a relationship while avoiding every conversation that could change it. One can read a medical report repeatedly while avoiding the decision for which the report was obtained.

Covert avoidance makes outcome measurement difficult. Attendance improves, yet participation remains narrow. The person completes the flight but spends the whole journey measuring the body and

rehearsing escape. A symptom scale may fall because the environment has been engineered carefully, even though flexibility has not increased.

This is why functioning cannot be inferred from symptom frequency alone. Fewer panic attacks may mean recovery, or it may mean that exercise, travel and solitude have disappeared. More activity may mean widening, or it may mean frantic control within a single safe corridor. Neither distress nor busyness tells the story without asking what became reachable.

An adequate account must also respect strategic avoidance. Not every confrontation is wise, and refusing an unsafe setting may be an expression of agency. The discriminating question is whether the action responds to present conditions and values, or whether an inherited alarm decides before the present situation has been sampled.

A policy, not a personality

Flow Hijacked uses the word policy for a pattern of action available to a system in a situation. It does not mean a conscious plan, and it does not turn a person into an algorithm. “Leave,” “check,” “ask,” “rehearse,” “delay” and “seek proximity” can all function as policies. Their value changes with state and context.

The Flow Hijacked synthesis proposes that repeated relief can concentrate choice around a narrow defensive repertoire. The empirical pieces-avoidance learning, relief, safety behaviour and costly persistence-do not by themselves prove this larger architecture. The synthesis becomes scientifically useful only if policy narrowing can be measured within people, predicts loss of valued action beyond symptom severity, and changes when corrective evidence becomes available.

Avoidance, then, is not the opposite of intelligence. It is intelligence operating on a very short invoice. It pays for the next minute with options borrowed from the next month.

And sometimes the body need not leave at all. A person can avoid uncertainty while sitting perfectly still, by thinking harder, checking inwardly, or asking another person to close a question that the world has left open.

Chapter 9: Worry, Reassurance, and the Search for Certainty

Thought as an emergency shift

At two in the morning, worry can feel like work. It has the posture of responsibility: turn the problem again, search for the missed angle, prepare for the sentence that might be spoken, rehearse the loss so it cannot arrive unannounced.

Some worry is useful. It can identify a solvable problem and produce an action: make the appointment, check the weather, apologize, revise the budget. Productive thinking usually changes form when it reaches a result. It yields a decision, a plan, or an explicit acceptance that no further action is available tonight.

Chronic worry often lacks that stopping rule. It keeps the mind close to the threat without necessarily bringing it closer to a decision. The person is active but not arriving. Questions multiply faster than answers because every answer can be evaluated for what it failed to guarantee.

“The doctor said the result was normal.”

Yes, but was the test sensitive enough?

“Your friend said she was not angry.”

Yes, but perhaps she was being polite.

“You prepared carefully.”

Yes, but what if the unexpected question is the one that matters?

The search is not stupid. Its impossible specification is the problem. It has been asked to convert an open future into a closed file.

What intolerance of uncertainty does and does not mean

Intolerance of uncertainty is an unfortunate name if it sounds like a character flaw. It refers to a tendency to experience uncertain situations as especially difficult or costly, cognitively, emotionally or

behaviourally. It is not the inability to understand probability. Nor is it unique to generalized anxiety.

A preregistered meta-analytic structural model published in 2026 synthesized 738 effects from 115 studies or theses, 138 independent samples and 28,693 participants. It found a strong association between intolerance of uncertainty and worry, approximately $r = .59$. The association weakened when depression and overlap among self-report measures were considered. ^[UW-01] Earlier meta-analytic work also connected intolerance of uncertainty with symptoms beyond generalized anxiety, including depression and obsessive-compulsive phenomena. ^[UW-02]

The empirical finding is an association across largely questionnaire-based literatures, not proof that one hidden “uncertainty mechanism” causes every form of worry. Shared wording can inflate relationships. People differ. Diagnoses overlap. A score cannot tell us which uncertainty matters to one person or what action follows from it.

The broader interpretation is nevertheless useful: uncertainty can carry a cost of its own. Two futures with the same average outcome may not feel equal if one is known and the other unresolved. A mediocre certainty can defeat a potentially better possibility because uncertainty itself is treated as part of the loss.

Reassurance is a social action

Reassurance looks like information, but it is also an event between nervous systems. A child asks whether the parent will return. A patient asks whether the sensation is dangerous. A partner asks whether the relationship is all right. The answer may supply facts, warmth, orientation and genuine co-regulation.

The trouble begins when reassurance must do something no relationship can do: certify the future.

Relief after an answer can teach the same lesson as relief after leaving a lift. Ask, receive certainty, settle. When doubt returns, asking becomes more compelling. Over time the question may become less about the external issue and more about whether the internal

feeling of certainty has arrived. “Do you think I am okay?” becomes “Are you sure?” and then “Are you sure you are sure?”

This does not make reassurance bad. Medical clarification, honest feedback and emotional support are part of sane life. The functional question is whether an answer supports action or becomes a ritual prerequisite for it. Good information eventually gives the person something to do or a boundary around what cannot yet be known. Compulsive reassurance asks the answer to remove the state of not knowing itself.

Worry as protection from surprise

One influential account proposes that worry can maintain a moderately negative emotional state and thereby reduce the contrast if something worse occurs. Experimental, prospective and momentary-assessment studies provide partial support for this contrast-avoidance model. ^[UW-04] ^[UW-05] ^[UW-07] But it is not a universal explanation of worry, and validation of a questionnaire is not validation of an entire mechanism. ^[UW-06]

The idea is psychologically recognizable. Hope can feel dangerous because disappointment falls farther from it. Relaxation can feel reckless because bad news would arrive without armour. Worry becomes a way of paying a little distress continuously to avoid being emotionally ambushed.

Yet this insurance policy changes the thing it insures. The feared event may never occur, but the preparation has already consumed the evening. A person who worries before every social encounter may attribute any later relief to having worried enough. Once again, protection competes for causal credit.

This is interpretation, not a verdict on why any particular person worries. Worry may also function as problem solving, verbal avoidance of imagery, learned responsibility, habit, self-criticism, or an attempt to preserve attachment. Several functions can coexist.

When worry becomes a moral duty

Worry is often protected by ethics. “If I stop thinking about it, I am careless.” “A good parent anticipates every danger.” “If I relax while someone I love is ill, I have abandoned them.” In such cases, asking the mind to worry less can sound like asking the person to care less.

The confusion is understandable because concern and worry share an object. Both may be directed toward a child, a relationship or a fragile future. Their consequences differ. Concern can organize attention and action, then make room for eating, sleeping and contact. Repetitive worry may consume those capacities while producing no new protection.

The useful distinction is not caring versus indifference. It is whether the mental activity changes a decision, prepares a viable response, or merely proves loyalty through suffering. Pain is not a receipt that demonstrates love.

This interpretation should be offered gently. For some people, vigilance was once required, and lowering it had real costs. Family or culture may also praise anticipatory responsibility. The aim is not to strip meaning from care, but to ask whether care has become fused with a process that leaves the cared-for person, and the carer, with a smaller life.

Information seeking and the moving finish line

The internet can answer more questions than any previous environment and can therefore generate more reasons for one further search. A symptom search begins with a common explanation, passes through a rare disorder, and ends with the discovery that no list is complete. Information increased; uncertainty did not reliably decline.

The problem is not “too much information” in the abstract. It is a stopping rule that moves whenever it is approached. Search until there is enough evidence for a reasonable decision is finite. Search until no threatening possibility remains is structurally endless.

The Flow Hijacked synthesis proposes that worry, reassurance and checking can function as forms of avoidance when they repeatedly close uncertainty without allowing valued action under uncertainty. This is not an established unitary mechanism. It should be tested by asking whether the behaviour reduces distress immediately, whether it narrows subsequent sampling, and whether the person's accessible actions depend increasingly on obtaining certainty first.

Sometimes the search turns inward. The next object to be checked is not the message, scan or relationship. It is the body itself.

Chapter 10: The Body Before the Explanation

The room changes before the story arrives

You enter a meeting and feel wrong before you know why. The shoulders rise. Breathing moves higher in the chest. Attention sharpens around faces. Only then does thought offer an explanation: they are judging me; something is about to happen; I should leave.

We often tell the story in the opposite order, as though a sentence appeared in the mind and pressed a button marked BODY. Human regulation is less tidy. Sensory cues, memory, posture, sleep, hormones, pain, social context and expectation interact continuously. A bodily shift can begin before a reportable interpretation, and the interpretation can then intensify the bodily shift.

“Before” does not mean outside the brain, beneath intelligence, or more truthful. It means that conscious explanation is not always the first event we can observe.

This matters because people can feel accused by cognitive accounts of anxiety. If no deliberate thought preceded the racing heart, being told to correct the thought may sound like being told the experience did not happen. The body did happen. The useful question is what the change predicts, what meaning it acquires, and what action it recruits.

A body prepared for action

Anxiety can alter breathing, muscle tone, heart rate, gastrointestinal activity, temperature, attention and the felt scale of time. No single pattern appears in everyone, and many changes are shared with exertion, excitement, illness, medication effects, sleep loss and ordinary stress. Physiology supplies context-sensitive resources; it does not stamp an emotion's name onto the skin.

This is why bodily arousal is both informative and ambiguous. A faster heart may accompany danger, but it may also accompany stairs,

coffee, fever, desire, dehydration or anticipation. The signal is real. Its cause and significance are not delivered with it.

Large literatures have searched for bodily signatures of anxiety. Some group-level differences exist—for example, lower heart-rate variability has been reported on average in anxiety-disorder groups—but heterogeneity, medication, comorbidity and measurement choices limit any individual inference. ^[NB-09] It would be wrong to translate a group difference into “your vagus nerve is deficient,” just as it would be wrong to infer safety from one normal physiological number.

The empirical finding is therefore narrower than popular diagrams suggest: anxiety involves bodily regulation, but there is no single bodily meter that diagnoses it or explains every episode.

The body as evidence

Imagine waking with a heavy chest after a poor night. The sensation changes the prior plausibility of the day. A neutral email feels more ominous. A crowded train seems less manageable. A choice that looked tolerable yesterday now appears reckless.

Nothing here requires the body to “cause” the thought in a simple one-way chain. The body changes the state from which possibilities are evaluated. Thought changes what the body is asked to prepare for. Attention selects further signals. Action changes the environment. The loop has no single royal entrance.

The interpretation offered in this book is that bodily state can modify the competition among predicted futures. When arousal is high, futures requiring endurance, ambiguity or social exposure may lose value; immediate protection may become more attractive. This is a broader scientific inference from multiple traditions, not a finding demonstrated by one decisive experiment.

Flow Hijacked goes one step further and labels this a possible component of policy narrowing. That is an explicit synthesis. It predicts that the same person will show a smaller range of reachable, valued actions at some bodily states than at others, even when external opportunity remains similar. It would be weakened if careful

within-person measurement showed no such state dependence beyond fatigue, illness, resources or ordinary preference.

Monitoring changes the monitored system

Place attention on the tongue for ten seconds. It becomes difficult not to notice saliva, pressure and the strange problem of where the tongue normally rests. The body did not suddenly acquire a tongue. Attention altered access to its signals.

Bodily monitoring can work similarly. Checking the pulse increases the number of beats available for interpretation. Watching breathing makes ordinary irregularities conspicuous. Scanning for dizziness can turn small shifts in balance into events. Attention does not invent every sensation, but neither is it a passive camera.

Monitoring can also alter physiology. A person who detects a slightly forceful beat may hold the breath, stiffen and release stress hormones; the next beat is now genuinely different. The observation has entered the system it observes.

This is not evidence that symptoms are “all in the mind.” It is evidence that minds, bodies and actions are not separate departments. A frightening interpretation can make a real sensation stronger; a real sensation can make a frightening interpretation more plausible.

Respecting medical uncertainty

There is a dangerous shortcut in both directions. One is to treat every benign bodily variation as evidence of catastrophe. The other is to dismiss every symptom as anxiety. New, severe, rapidly changing or medically concerning symptoms require appropriate assessment. A book cannot examine a body, establish a diagnosis or decide which risk is reasonable for one reader.

The point is not that chest pressure, dizziness or breathlessness is harmless. The point is that sensation alone does not identify its cause. Responsible interpretation uses context, history, medical information and change over time.

Once urgent causes have been appropriately considered, a second question becomes available: what relationship has formed between sensation and prediction? That question belongs to interoception-the nervous system's ongoing encounter with the body's internal world.

Bodies borrow context from other bodies

Physiology is also social. A calm voice can slow the pace at which ambiguous sensations are interpreted; an alarmed face can make the same sensations more consequential. This is not magic transmission and not proof that another person can regulate us from the outside like a thermostat. It is ordinary evidence: posture, timing, expression, touch, distance and language alter what the situation appears to require.

The effect is not uniformly soothing. A familiar person may signal safety in one history and danger in another. Reassurance can open exploration or become its required condition. The scientific finding is therefore never simply “company calms the nervous system.” Person, relationship, age, task and context matter.

The broader interpretation is that bodily state cannot be understood as a sealed property inside one skin. Social surroundings enter prediction and action. Flow Hijacked will later ask whether support expands the actions a person can take or merely suppresses distress while the same boundary remains. Those outcomes must not be confused.

Chapter 11: Interoception, When Sensation Becomes Evidence

More than hearing the heartbeat

Interoception is often described as the sense of the body's interior. The phrase is useful, but it can make the body sound like a control panel that sends clean readings to a spectator in the head. What arrives is messier: pressure, temperature, stretch, chemical state, breathing, heartbeat, hunger, nausea, pain, effort and many signals that never become conscious at all.

The nervous system does not merely receive these signals. It regulates the organs producing them, predicts what should be arriving and decides which deviations deserve action. Interoception is therefore closer to a conversation than a dashboard.

Several distinctions prevent that conversation from collapsing into one score. Interoceptive accuracy concerns performance on a task—whether a person can discriminate or detect a bodily event. Interoceptive sensibility concerns beliefs and reports about bodily attention or sensitivity. Interoceptive awareness can refer to insight into one's own accuracy. Attention asks how often signals are sampled. Interpretation asks what they mean. A person may notice many sensations, feel highly confident about them, interpret them catastrophically and still perform no better on an objective detection task.

This distinction corrects a popular mistake. A systematic review and meta-analysis of 55 studies found no overall association between anxiety and objective cardiac interoceptive accuracy. ^[IP-01] The heartbeat tasks themselves have important methodological problems, so this is not proof that bodily processing is irrelevant. It does mean we should stop saying, without qualification, that anxious people simply feel their hearts more accurately.

Precision is not the same as truth

In ordinary language, a precise person is often a correct person. In predictive accounts, precision has a different meaning: roughly, the weight or confidence assigned to a signal or expectation. A sensation can carry great weight without identifying its cause correctly. A forecast can dominate action without being probable.

Suppose a person notices three forceful heartbeats. The raw event is already partly constructed—attention has selected it from thousands of bodily fluctuations. The event then enters a field of possible explanations: exertion, caffeine, poor sleep, emotion, illness, danger. History, context and current state change the weight assigned to each.

If the dangerous explanation gains precision, attention returns to the chest. More beats become noticeable. Breathing changes. The new sensations are not fabricated; they were partly generated by the attempt to determine whether the first sensations were dangerous. Evidence and inquiry become entangled.

The scientific interpretation is not that the brain hallucinates the body at will. Sensory constraints matter, disease matters, and prediction can be wrong in both reassuring and alarming directions. The narrower point is that bodily meaning is inferred under uncertainty.

Anxiety sensitivity: fear of what fear might do

Two people can experience similar arousal and inhabit different events. One reads a pounding heart as effort. Another reads it as the beginning of collapse. The difference is not merely intensity; it is expected consequence.

Anxiety sensitivity names the tendency to fear anxiety-related sensations because of beliefs about their physical, psychological or social consequences. Meta-analytic evidence shows substantial associations with anxiety disorders, especially panic, but also beyond it. ^[IP-04] Much of this literature uses self-report, and association does not establish a single cause. Anxiety sensitivity is neither a diagnostic

test nor an explanation that replaces learning, context or medical assessment.

Still, the construct illuminates a peculiar recursion. A person is not only afraid. They are afraid that fear will damage the heart, expose weakness, produce insanity, cause fainting or become unstoppable. The alarm becomes an object of alarm.

That second-order fear changes policy. Exercise is avoided because breathlessness resembles panic. Warm rooms are avoided because heat resembles illness. Excitement becomes suspicious because it raises the pulse. The range of bodily states permitted inside ordinary life begins to contract.

Confidence about sensing

There is another underappreciated possibility: the difficulty may lie not in sensing too much or too little, but in knowing when to trust a sensation. Low confidence invites repeated checking. Excessive confidence in a catastrophic reading can make alternatives inaccessible. Both can produce more sampling.

Consider a person counting their pulse at the wrist. A skipped count may become evidence of an irregular beat; uncertainty about the count prompts another trial; the repeated trial increases attention and arousal. Soon the person is investigating not only the heart but the reliability of their own investigation.

This metacognitive layer matters because reassurance at the first level may not resolve doubt at the second. A normal reading answers “What did the device show?” It does not necessarily answer “Can I trust the device, the timing, my interpretation, or my ability to know?” Certainty retreats one level upward.

Associative experiments have examined links between bodily sensations and catastrophic interpretations, but such tasks are indirect and do not establish that a measured association caused natural panic attacks. ^[IP-06] The finding should sharpen questions, not close them.

The body as a teacher, not a verdict

Interoceptive exposure in cognitive-behavioural treatment deliberately evokes selected sensations-through safe, clinically appropriate exercises-so that predictions about them can be tested. Component meta-analytic evidence suggests that packages including interoceptive exposure contribute to effective treatment for panic, although components are difficult to isolate and comparisons include indirect evidence. ^[IP-05]

The principle is larger than any one exercise. A body can supply corrective information only if variation is allowed and meaning remains revisable. If every rise in pulse is stopped instantly, the system receives abundant evidence about successful stopping and little evidence about the natural arc of arousal.

The Flow Hijacked synthesis proposes that interoceptive fear can reduce future reachability by excluding activities, places and relationships that might generate disallowed sensations. This is not established by the interoception studies alone. It is a hypothesis linking bodily meaning to action diversity. It would require within-person evidence that changes in feared bodily states predict changes in valued action, not simply changes in symptom ratings.

Panic is what happens when this conversation accelerates until the body seems to deliver a verdict before any alternative can be heard.

Chapter 12: Panic, Catastrophe From Within

When seconds become an emergency

A first panic attack can divide biography into before and after. The person may be on a train, in a supermarket, at a desk, or apparently doing nothing. The heart surges. Breathing feels wrong. Vision narrows. The body becomes unfamiliar at the exact moment it becomes impossible to ignore.

Then comes the catastrophic question: Am I dying? Am I fainting? Am I losing control? The question is not leisurely. It arrives with the felt authority of the sensations it is trying to explain.

Panic compresses time. A medical appointment next week is irrelevant to a body demanding an answer in the next ten seconds. A reassurance that symptoms usually pass may be factually sound and temporally useless. The system is selecting an emergency action now.

Escape, sitting down, gripping a surface, calling someone, checking oxygen, searching symptoms—each action may reduce danger, distress or uncertainty. Each can also become associated with survival. Afterward, the person faces a new threat that did not exist before the attack: the possibility of another attack.

A loop without a single starting point

The classic panic loop is easy to draw. Sensation is noticed. Sensation is interpreted as dangerous. Fear increases sensation. Increasing sensation strengthens the interpretation. But real episodes do not always begin in the same corner. Poor sleep may alter breathing. A memory may shift the body. A crowded room may increase vigilance. A spontaneous sensation may arrive first. Sometimes the person notices only after the loop is already moving.

Anxiety sensitivity is strongly associated with panic, but it does not explain every attack or establish a universal causal route. ^[IP-04] Carbon-dioxide provocation studies show that induced panic is more likely in

several diagnostic groups than in controls. A network meta-analysis included 41 studies and 2,181 participants. ^[IP-02] Yet carbon-dioxide sensitivity is not completely specific to panic disorder, and a provocation in a laboratory is not a miniature copy of a spontaneous attack in a life.

These findings support bodily reactivity as one part of the landscape. They do not identify a panic molecule or an inevitable personal mechanism.

The amygdala is not a panic switch

Popular neuroscience often gives panic a glowing dot. The amygdala becomes the alarm, and the rest of the person becomes its audience. This is too simple.

In a striking lesion report, three people with rare bilateral amygdala damage experienced panic during carbon-dioxide inhalation. ^[IP-03] The sample was only three, and the condition was exceptional. The study does not show that the amygdala is unimportant to external threat. It does provide a useful counterexample to the claim that an intact amygdala is necessary for every human panic response.

Imaging meta-analyses instead implicate distributed systems involving salience and interoceptive regions, cingulate and prefrontal areas, memory and context systems, amygdala, brainstem and periaqueductal grey, with patterns varying by task and kind of threat. ^[NB-02] ^[NB-03] Brain imaging is correlational and group-based; it cannot tell a reader which circuit “caused” an individual episode.

The better picture is not a command centre but a coalition. Body, prediction, attention, memory and action can align rapidly enough to feel like one indivisible fact.

From panic to the fear of places

After an attack, the person often reviews geography. Where could I escape? Where would help be available? What if this happened on a motorway, in a queue, during a haircut, halfway through a flight?

Places acquire meaning through the actions they permit. A cinema seat in the middle of a row is not threatening because of upholstery. It is threatening because exit would be visible and slow. A bridge is not only high; it limits alternatives. Being alone is not merely solitude; it changes access to rescue or reassurance.

This is how panic can recruit agoraphobic restriction. The feared object is not necessarily the place. It may be the combination of sensation, anticipated catastrophe and limited escape. The world is reorganized according to response options.

Avoidance then protects against panic and against learning about panic. Fitness may fall, making exertion more intense. Safe zones become more exact. A companion's presence may permit travel, but safety may be credited exclusively to the companion. The first episode becomes an event; the anticipation of recurrence becomes an environment.

When vigilance prolongs recovery

Even relief can become suspect. After the peak, a person may monitor whether the body is truly settling. Is the heart normal yet? What if the attack returns? Monitoring recruits attention, and attention may keep small fluctuations prominent. Recovery becomes another test that can be failed.

This creates a clinically important asymmetry. The threshold for entering alarm may be higher than the threshold required to remain there. Once the system is mobilized, residual sensations, memory and post-event checking can sustain it. Later we will call this possible path dependence hysteresis. For now, it is enough to notice that the way into a state and the way out need not be mirror images.

The Flow Hijacked synthesis proposes that panic can narrow policy before the next attack occurs. Anticipated bodily catastrophe assigns veto power to situations that might generate or trap arousal. That is a model-level interpretation, not a validated panic biomarker. It should be tested against simpler explanations such as symptom severity, physical capacity, depression and actual environmental constraints.

Panic is not cowardice and not theatricality. It is a state in which internally generated evidence acquires emergency weight. Understanding that weight requires us to ask what prediction is for in the first place.

PART IV

THE PREDICTIVE FUTURE

Uncertainty, priors, precision, learning, and the machinery of tomorrow.

Chapter 13: The Threat-Prediction Machine

Prediction is not prophecy

Calling the nervous system predictive can sound grander than the claim needs to be. It does not mean the brain sees tomorrow. It means present regulation depends partly on estimates of what is likely to happen next.

Reach for a glass and the hand is shaped before contact. Step off a kerb and posture changes in anticipation of the drop. Hear the beginning of a familiar sentence and some continuations become more available than others. Prediction is woven into perception and action because waiting for completed events would make coordinated behaviour impossible.

Threat prediction is the same ordinary capacity under expensive conditions. If an organism waits until injury is certain, defence is late. It must act using incomplete information, balancing false alarms against missed dangers. Human beings add symbolic futures: rejection, diagnosis, guilt, unemployment, disgrace, abandonment. An event need not be physically present to alter the body preparing for it.

Prediction therefore has a practical job. It allocates attention, energy and action before certainty. The central problem of anxiety is not that prediction occurs. It is how predictions are weighted, updated and allowed to govern.

No single fear room

The attraction of a “fear centre” is obvious. It turns a distributed problem into a switch. But meta-analytic neuroimaging does not support one location that cleanly contains anxiety.

Across 181 studies, induced and pathological anxiety showed overlapping activity particularly in cingulate and insular regions within a wider distributed pattern. ^[NB-02] Meta-analysis of human fear-conditioning studies likewise implicated networks involved in salience,

interoception, autonomic regulation and action, not one isolated centre. ^[NB-03] These are group-level, task-dependent associations. They do not provide a diagnostic scan or establish that activation in a region is the cause of a person's fear.

Threat imminence also matters. In a simulated-predator task, more distant threat emphasized more anterior systems, whereas near threat was associated more strongly with midbrain and periaqueductal-grey responses. ^[NB-05] The study was small and specific to its task; it does not prove a universal hierarchy. It does illustrate a broader principle: a threat ten minutes away and a threat one second away pose different control problems.

Temporally uncertain anticipation has been associated with a broad network including midcingulate and insular cortex, frontal areas, amygdala/BNST-related regions, periaqueductal grey and brainstem. ^[NB-07] Again, activation does not equal mechanism. The empirical picture is distributed and conditional.

The forecast enters the present

Suppose tomorrow contains a difficult conversation. The forecast changes sleep. Poor sleep changes irritability and bodily arousal. Arousal changes how an ambiguous expression is read. That reading changes what is said. What is said changes the conversation that actually occurs.

Prediction has not reached backward through time. A present representation of the future has changed the path into that future.

This recursive structure explains why anxious predictions can feel self-confirming without being delusions. “The meeting will go badly” produces guarded speech. Guarded speech makes connection harder. A colder interaction follows. The outcome contains both external reality and the policy recruited by the forecast.

The scientific interpretation is that anxious anticipation should be studied as an action-perception loop, not as a private sentence waiting to be fact-checked. This does not erase distorted beliefs or genuine hazards. It places belief inside a moving system.

Prediction errors that fail to teach

In a simple learning story, the system predicts an outcome, observes what happens and updates according to the mismatch. Real life complicates every term.

What exactly was predicted? That I would blush, or that everyone would despise me? What was observed? A neutral face, an ambiguous silence, a kind remark dismissed as politeness? Who received causal credit for safety-the world, my ability to cope, or the ritual I performed?

An apparently disconfirming event may therefore produce little revision. The prediction was vague enough to migrate. The outcome was ambiguous. Protection explained the success. Attention missed the safe information. Or the person updated intellectually while action values remained organized by relief.

This is not evidence that beliefs are sealed. It is a reminder that updating depends on representation, sampling and attribution. “Nothing bad happened” is not a complete account of what the system had a chance to learn.

Predictions are negotiated with other people

Many important forecasts concern minds we cannot inspect. Does this person still trust me? Did that pause mean disappointment? Will my uncertainty change how I am seen? Social prediction is difficult because its evidence is partly produced through interaction.

A guarded person receives different responses from an open one. Repeated reassurance can soothe a relationship and also strain it. Silence may protect against embarrassment while depriving the other person of information needed to respond warmly. The forecast enters the exchange, and the exchange returns as evidence.

None of this means that rejection is self-created. Other people have their own histories, motives and power, and social danger can be entirely real. The interpretation is narrower: interpersonal evidence is

not collected from a fixed environment. It emerges between agents who change one another.

This makes social learning especially vulnerable to causal ambiguity. A kind response may be credited to perfect self-monitoring; a cold response may be credited to personal defect rather than to the other person's state; no response may support whichever story was already most available. A serious model must leave room for both actual relationships and the policies through which they are sampled.

Finding, interpretation, synthesis

The empirical findings in this chapter concern distributed, task-dependent recruitment during conditioning and anxious anticipation.

[NB-02] [NB-03] [NB-07] They do not demonstrate a unified predictive-processing theory of anxiety.

The broader interpretation is that threat forecasts alter current regulation and evidence sampling. This is consistent with learning, interoception and action research, but several formal theories can express it.

The Flow Hijacked synthesis proposes a more specific architecture: when one threatening future acquires enough weight, it can reorganize present body, attention and policy so that competing futures become less reachable. That proposal will later require operational definitions and tests. At this stage it is a disciplined question, not a result.

To understand what a future representation can do, we next need to understand how the mind constructs a place it has never visited.

Chapter 14: The Future as a Place of Danger

Tomorrow is built from fragments of yesterday

Imagine breakfast next Thursday. A table appears, perhaps light through a window, perhaps the taste of coffee. None of this is a stored recording of an event that has not happened. The mind assembles a plausible scene from memory, knowledge, habit and current concern.

Prospection is constructive. That is why it is powerful and fallible. We can rehearse a conversation, anticipate obstacles and compare paths without physically trying them. We can also mistake vividness for likelihood and familiarity for inevitability.

An anxious future is often not an empty future. It can be extraordinarily detailed. The problem is distribution. Threatening branches receive repeated simulation; ordinary and desirable branches remain sketchy. The catastrophe is rendered in high resolution while coping appears as a blank after the words “and then.”

This asymmetry changes present value. A vivid event is easier to feel now. The body prepares for an image whose sensory detail gives it the force of proximity. Yet vividness is not evidence of probability. It is a property of representation.

The research literature on episodic future thinking is broad and mixed, and the current evidence ledger does not justify a single strong claim that anxious people universally imagine the future less specifically or more vividly. The responsible position is narrower: future simulation varies across people and states, and its relation to anxiety depends on valence, specificity, depression, task and measurement. Here the book uses prospection as an explanatory lens, not an individual diagnostic fact.

The missing chapter after catastrophe

Ask someone what they fear about a medical result and they may describe the telephone call, the doctor's pause, the faces of family

members. Ask what happens after the feared moment and the simulation can stop.

But lives do not stop where anxious images stop. After bad news come decisions, support, adaptation, second opinions, treatment, grief, boredom, humour, meals and mornings. This is not an argument that outcomes are secretly good. It is an observation that catastrophic simulation often models impact more richly than recovery.

The omitted variable is coping. Probability asks, "Will this happen?" Cost asks, "How bad would it be?" Coping asks, "What could I do next, with whom, and what capacities might become available then?" Anxious reasoning can estimate the first two repeatedly while leaving the third almost unrepresented.

Cheap reassurance says, "You will be fine." A fuller future says, "The outcome is not guaranteed, and the story contains more steps than the feared scene."

The veto branch

Consider a journey with ten plausible outcomes. Nine are ordinary or rewarding. One includes panic. If the panic branch carries enough subjective cost, it can veto the journey even when no one believes it is the most likely branch.

This differs from a future that has lost all attraction. In some depressive states, positive futures may lose pull or availability. In anxiety, a valued future can remain vivid and desirable yet inaccessible because one threatening branch has acquired stopping power. The person wants the journey and does not take it.

This is an interpretive distinction, not a clean diagnostic boundary. Anxiety and depression frequently coexist, and both can affect future thinking. The point is functional: a future can fail to guide action because it has no pull, because threat vetoes it, or because external conditions make it impossible. These are not the same problem.

Imagination as rehearsal and exposure

Simulation can protect. Pilots rehearse emergencies; surgeons plan complications; families prepare for storms. Useful rehearsal specifies an action and ends when preparation is adequate. Anxious simulation can become rehearsal without action, repeated not to prepare but to feel briefly less unprepared.

The difference is visible in what changes. After useful rehearsal, a plan improves, a resource is gathered, or a decision is made. After circular simulation, the scene may simply acquire more detail and emotional authority.

Yet imagination is not only a maintenance mechanism. It can also extend reach. A person can simulate coping, variability and continuation. They can imagine speaking with a trembling voice rather than only speaking without fear. They can include uncertainty in the scene rather than waiting to imagine a guaranteed outcome.

The Flow Hijacked synthesis suggests that future accessibility should be measured not by the number of scenarios a person can name but by the number of valued scenarios that remain behaviourally enterable under uncertainty. This is a proposed construct. A diary full of plans may coexist with one narrow policy; a quiet week may contain genuine openness.

The obstacle is often not a lack of futures. It is the price charged for entering one before its ending is known.

Memory is a library with a changing catalogue

Future construction depends on memory, but memory is not a neutral warehouse. What becomes easy to retrieve depends on mood, context, recency, repetition and current goals. A frightened state makes some episodes unusually available: the flight that shook, the doctor who once missed something, the conversation that ended badly.

Availability can masquerade as frequency. If three threatening examples arrive quickly and safe examples require effort, the threat

may feel representative even when it is rare. The error is understandable; minds cannot search an entire lifetime before every decision.

The implication is not that people should replace negative memories with positive slogans. A useful future needs credible material. It may include previous recovery, support that was actually available, skills acquired since the event, and the ordinary fact that outcomes vary. The aim is not optimism. It is a less selectively catalogued imagination.

This is an interpretation of how memory can bias prospection, not a claim that retrieval alone explains anxiety. Current danger, learning, depression and social conditions can all determine which futures are realistic.

Chapter 15: Uncertainty Is Not Neutral

The waiting room with no clock

There are difficult events, and there is waiting to learn which event has occurred. The two are not psychologically interchangeable.

A train delayed by exactly forty minutes may be inconvenient. A train delayed “until further notice” can reorganize the whole station. The expected delay may even be shorter, yet the inability to plan becomes a cost in itself. Should I leave the platform? Call someone? Buy food? Keep watching the board?

Uncertainty withholds the condition for committing to an action. That is why it can feel active rather than empty. The system keeps policies available, samples for updates and resists investing fully in any one path.

For some people, this suspended state becomes more aversive than a moderately bad certainty. A familiar job that constricts life can defeat an uncertain opportunity. Predictable loneliness can feel safer than uncertain intimacy. The choice is not necessarily a mistaken estimate of average reward. The variance, ambiguity and perceived inability to recover also enter the price.

One phrase, several experiences

“I don't know” can refer to different problems.

We may not know **what** will happen: whether the result is positive or negative. We may not know **when** it will happen: today or next month. We may not know **how variable** the world is: a bad day or a changed relationship. We may not know **whether our information is adequate**: an unread message, an imperfect test. We may know the probabilities and still not know **whether we could cope**.

These uncertainties recruit different actions. Timing uncertainty promotes monitoring. Ambiguity promotes interpretation. Doubt about information promotes checking. Doubt about coping promotes

proximity and escape planning. Collapsing all of them into one intolerance score can hide the mechanism that matters.

The 2026 meta-analysis connecting intolerance of uncertainty and worry provides strong corpus-level evidence of association, not a map of these distinct routes and not causal proof. ^[UW-01] The overlap with depression and with questionnaire content is part of the result, not an inconvenience to be deleted.

The emotional cost of surprise

One reason uncertainty can be expensive is that it permits emotional contrast. A relaxed state may be followed by bad news. A hopeful state may be followed by disappointment. Worry can keep mood closer to the feared destination, reducing the drop if something bad occurs.

The contrast-avoidance account has empirical support from experimental, prospective and everyday assessment studies, but none establishes it as the sole engine of generalized anxiety. ^[UW-04] ^[UW-05] ^[UW-07] A person may worry for several reasons, and a statistical association cannot identify the function of one episode.

Still, the account captures something humane: sometimes people do not fear only pain. They fear being caught without the right emotional posture. “If I relax, it will be worse.” “If I hope, I will have tempted the fall.” Worry becomes advance payment on grief.

The payment is real even when the bill never comes.

Uncertainty in a genuinely dangerous world

Any serious theory of anxiety must survive contact with real uncertainty. People wait for biopsy results, live through war, navigate unstable work, care for ill relatives and inhabit relationships in which trust has been broken. The aim cannot be to persuade them that the world is safe.

Adaptive vigilance responds to conditions. It changes when information changes, distinguishes actionable from unactionable risk and preserves resources for more than one possible future. A locked

system may continue the same policy across contexts, timescales and evidence.

This is not a moral distinction between brave and anxious people. Resources, history, discrimination, disability and actual exposure change what counts as a viable action. Sometimes the narrow route is the only safe route. A model that calls every protection pathological is itself poorly calibrated.

The interpretive task is to ask whether uncertainty is being priced according to the situation or has become a general surcharge on living. Does the person gather information that changes a decision? Can monitoring stop? Can a valued action occur before certainty? Does the policy soften when danger falls?

From feeling to structure

Flow Hijacked proposes that high uncertainty cost may increase the immediate value of defensive control and reduce exploration. This is an explicit synthesis, not a direct finding of the questionnaire meta-analyses. To test it, uncertainty would need to be manipulated while observing real choices, corrective sampling and valued action within people-not merely asking whether uncertainty feels unpleasant.

The distinction matters. A person can dislike uncertainty and still move flexibly. Another can report moderate discomfort yet organize an entire life around preventing it. Experience and policy are related, not identical.

Before we can formalize any such proposal, uncertainty itself has to be disassembled. Some uncertainty belongs to the world. Some belongs to our knowledge. Some concerns changing rules. Some concerns whether we have the power to respond.

Uncertainty is not danger, but it can change danger's price

Uncertainty and threat are often spoken of as if they were synonyms. They are not. We can be certain that an event will be unpleasant, uncertain about a harmless surprise, or uncertain whether

a severe event will occur. The dimensions interact without collapsing into one another.

What uncertainty changes is the space of preparation. When timing, cause or available response is unclear, more policies must remain under consideration. That can be metabolically, attentively and emotionally expensive. The cost is especially high when the person believes one missed signal would be unforgivable.

This distinction protects against a common exaggeration. Anxiety is not produced by uncertainty alone, and reducing uncertainty is not always possible or desirable. The relevant question is how uncertainty changes the weighting of threat, control and action in a particular context.

Chapter 16: The Shapes of Uncertainty

Dice, fog and moving rules

Imagine three games.

In the first, a fair die is rolled. You know the possible outcomes and their probabilities, but not which one will occur. The uncertainty is built into the event.

In the second, a die is rolled behind a screen, and you do not know whether it is fair. More information could improve the model.

In the third, the rules governing the die change without announcement. Evidence collected yesterday may no longer describe today.

These are not the same ignorance. They roughly distinguish randomness in outcomes, uncertainty about the model, and volatility in the process generating outcomes. Technical traditions use terms such as aleatoric and epistemic uncertainty, stochasticity and volatility, with definitions that vary by field. The names matter less than the control problem: should I accept irreducible variation, seek information, or detect that the world has changed?

An anxious system can make errors in any direction. It may search for information where no amount can remove chance. It may accept a poor model as settled. It may interpret ordinary noise as evidence that the rules changed. Or it may fail to update when change is real.

Ambiguity is not probability with the numbers missing

A probability can be calculated only after a model has specified relevant possibilities and relations among them. Many human questions do not arrive in that form. Does this silence mean anger, fatigue, distraction or nothing at all? Will a treatment help this individual? Is a new career worth leaving the old one?

Numbers can discipline reasoning, but they cannot rescue an underspecified model by themselves. The same observation can

support several causal stories. A friend replies after reassurance: perhaps reassurance repaired a misunderstanding, perhaps the friend was always available, perhaps repeated asking created strain that has not yet appeared.

This is why “think statistically” is good advice and an incomplete therapy. Base rates matter. So do model choice, evidence quality, cost, agency and the opportunity to learn.

Volatility and the problem of when to forget

A stable world rewards slow updating. One cold email should not overturn years of trust. A changing world rewards faster updating. If a workplace has genuinely become unsafe, loyalty to the old model becomes costly.

The system must therefore estimate not only what is happening but how quickly its environment changes. Too much assumed volatility makes every fluctuation a possible regime shift. Too little makes the system cling to history after conditions have altered.

Research on temporally uncertain threat shows recruitment of distributed systems during anticipation, but it does not tell us that one person's distress is caused by a single volatility parameter. ^[NB-07] Computational constructs can organize experiments; they are not diagnoses in disguise.

The broader scientific inference is that anxiety may sometimes involve altered learning about change, but findings vary by task, population and model. We should resist turning an elegant parameter into a personality type.

Controllability changes the meaning of not knowing

Uncertainty about rain is different when you have a coat. Uncertainty about panic is different when you believe you can remain, recover and ask for help. The external probability may be unchanged; the decision value changes because action remains available.

This is why coping cannot be added as a cheerful afterthought. It is part of the formal problem. Expected loss depends on possible

outcomes and their costs, but those costs depend partly on what the organism expects it can do next.

Perceived control can also mislead. Repeated checking may feel controlling while adding no useful information. Conversely, accepting irreducible uncertainty may look passive while freeing resources for actions that genuinely alter the situation. Control is not the same as activity.

Exploration and exploitation

Decision theory uses a useful pair of words. Exploitation means choosing the option currently believed to work best. Exploration means accepting some uncertainty in order to learn whether another option may be better. Every adaptive system needs both. A life of pure exploration never consolidates knowledge; a life of pure exploitation becomes prisoner to its first adequate answer.

Threat changes the balance. When error is expensive, repeating the known safe policy is sensible. The difficulty arises when the policy prevents the evidence required to discover that conditions have changed. The safest known route keeps winning partly because alternative routes are never allowed to acquire value.

Exploration is not thrill seeking. It can be tiny: waiting five minutes before checking, asking one unrehearsed question, taking a familiar journey without one protective ritual. The size of the experiment should reflect actual risk, capacity and support. An experiment that overwhelms the system may teach only that exploration is dangerous.

The empirical avoidance studies establish pieces of this problem, not a complete computational account of a person's life. ^[LE-06] ^[LE-07] The Flow Hijacked synthesis predicts that anxiety becomes especially restrictive when exploitation of defensive policies suppresses informative exploration. A fair test would compare this account with simpler alternatives and ask whether carefully chosen exploration changes later policy rather than merely producing distress in the moment.

Measuring uncertainty without pretending to finish it

Questionnaires capture important subjective tendencies, and their large literatures show a robust relationship between intolerance of uncertainty and worry. ^[UW-01] They cannot alone tell us whether a person seeks information, avoids variance, updates too quickly, or refuses action until ambiguity disappears.

A fuller measurement strategy would separate at least four things: the uncertainty presented by the task; the person's report of how aversive it feels; the information they choose to sample; and the actions they remain willing to take. Repeating those measures over time would reveal whether the relevant pattern is stable, state-dependent or changed by context.

This is the point where interpretation can become a research programme. The Flow Hijacked synthesis predicts that uncertainty will matter most when it changes policy diversity and future reachability. Discomfort alone is insufficient. The hypothesis would be weakened if uncertainty ratings rose while valued action and information sampling remained flexible, or if apparent narrowing were fully explained by external constraint, depression, fatigue or lack of resources.

Formal language should make the model easier to fail, not harder to understand. We will later introduce equations only when each symbol names something that could, at least in principle, be observed. Before that, one more piece is needed: how prior expectations and incoming evidence share authority.

Chapter 17 - Prediction, Priors, and the Weight of Evidence

The chair you do not have to rediscover

You do not begin each morning as a visitor to an unexplained planet. You lower yourself into a chair without first proving that it will hold you. You hear the particular cough of a familiar engine and know who has arrived. You read a face before you could list the shadows, angles, and tiny muscular movements from which that impression was assembled.

The nervous system gets away with this astonishing efficiency by bringing yesterday into today. Perception is not a camera waiting politely for light. It is closer to a fast, informed negotiation between what arrives and what was already expected. A world encountered without expectations would not feel wonderfully open-minded. It would be unusable. Every cup would be an investigation; every footstep, an unsolved case.

In predictive language, an expectation carried into the present is often called a *prior*. The word can sound like courtroom furniture, but the idea is ordinary. Autumn clouds make rain more plausible than a flawless blue summer sky. A friend's usual five-minute delay does not mean what the same delay would mean from a surgeon during an operation. History supplies the starting point from which new evidence is interpreted.

Anxiety does not begin because the brain commits the elementary sin of expecting. Expectations are indispensable. Nor are strong expectations automatically a problem. If a road has been dangerous every night for a year, approaching it cautiously tonight is not a cognitive defect. The more useful question is whether an expectation remains answerable to the world in which it is now operating.

That question turns out to be much harder than “Is the thought rational?” A person may say, quite rationally, that catastrophe is unlikely while their attention, muscles, breathing, and available

actions are organized as though it has already won the election. Beliefs have more than one public face. What we can endorse in a quiet sentence and what controls us under pressure may be related without being identical.

A prior is history in travelling clothes

A prior is sometimes described as if it were merely a number stored in a mathematical cupboard. In a model it may be represented by a number. In a life, it arrives with a biography.

The expectation that dogs may bite can be shaped by a bite, by a parent's alarm, by a frightening story, or by years of watching other people cross the street. The expectation that a rapid heartbeat means danger can be condensed from one overwhelming panic episode. The expectation that mistakes invite humiliation may have been trained through a thousand small corrections delivered with contempt. None of these histories needs to be consciously recalled at the moment the expectation acts.

This matters morally as well as scientifically. A strong threat prior may have been learned in a world where threat was common. Hypervigilance in an unsafe home, concealment in a hostile social environment, or rapid compliance under coercion may be intelligent adaptations to the available evidence. Later, after the world changes, the rule may remain. The difficulty is not that the person failed to learn; it is that the old lesson may now be applied too broadly, too quickly, or without an effective way to discover that the conditions have altered.

Even then, caution about caution is warranted. There is no universal observer who can look at another person's life and announce that danger is over. Risk can be unevenly distributed. A room that is benign for one body may be dangerous for another. Calibration is therefore not conformity to the average person's confidence. It is the fit between a forecasting system, the actual conditions, the costs of error, and the person's capacity to respond.

The compressed history in a prior includes culture and institutions too. A patient repeatedly dismissed by clinicians does not approach the next appointment with the same starting assumptions as a patient repeatedly heard. Someone whose job has vanished twice during restructuring will read “quick chat tomorrow” differently from a colleague with a secure contract. Priors are personal, but they are not merely private inventions.

Precision is a volume control, not a truth detector

If several accounts of the same moment are possible, the system needs some way to decide which information deserves influence. Predictive models use the term *precision* for the estimated reliability or confidence assigned to a signal, prediction, or error. Precision determines volume in the negotiation. It does not certify truth.

A witness can be magnificently certain and wrong. A thermometer can display three decimal places while badly calibrated. An anxious prediction can carry great functional confidence even when the person explicitly calls it “probably anxiety.” Precision answers, in effect, “How much should this input move me?” It does not answer, “Is this input correct?”

Imagine a medical consultation. Tests have been reassuring. The clinician explains why a dangerous condition is unlikely. On the walk home, one forceful palpitation lands in the chest. The medical account is abstract, distributed across several facts, and now a few hours old. The palpitation is immediate, embodied, vivid, and accompanied by a burst of arousal. Which source gets the microphone?

There is no requirement that the newest or loudest source should win. Yet salience is easily mistaken for diagnostic value. A signal may receive high precision because it is intense, close, familiar, or emotionally expensive to ignore. Under threat, the system can shift the practical weighting of evidence long before the person has composed a new conscious conclusion. Research on temporally uncertain threat implicates a distributed set of cortical, subcortical, and brainstem systems rather than a single “anxiety centre,” and

those imaging results describe group-level task responses, not a scanner-based truth meter for an individual. ^[NB-07]

The same contest occurs outside the body. Ten colleagues respond warmly to a presentation; one looks bored. A partner offers months of affection; one message is delayed. A flight proceeds uneventfully; turbulence occupies the memory. The minority observation is not always irrelevant. Sometimes the bored colleague has seen a real problem. But threat can enjoy a peculiar evidentiary privilege: supporting details are treated as signals, while contradictory details are treated as noise.

The umbrella calculation nobody writes down

You wake beneath a grey sky. Yesterday's forecast mentioned rain. You take an umbrella. No one applauds the Bayesian sophistication of this decision, which is fortunate, because you have not performed explicit arithmetic. You have combined prior experience, current evidence, and the cost of being wrong.

Now alter one fact. Last week an unexpected storm destroyed your laptop. The same clouds may produce a much stronger impulse to prepare. The probability of rain need not have changed. The cost landscape has.

This is why disagreements about “how likely is it?” often miss the mechanism that controls action. Two people can agree on likelihood and choose differently because one predicts a much greater loss, carries less confidence in recovery, or places a larger moral penalty on having failed to prevent the outcome. Forecasting and valuing are braided together in behaviour even when a theorist separates them on paper.

The umbrella also reveals something subtler. Carrying it changes the meaning of the afternoon. If rain arrives, the umbrella was wise. If rain does not arrive, it did no great harm. Defensive actions in anxiety can have a similar asymmetry, except that some are costly and some obscure the evidence. If a person enters a meeting only after rehearsing every sentence, a successful meeting can be credited to

the rehearsal. If the preparation was exhausting, that cost may still feel preferable to discovering whether less preparation would have worked.

A prior therefore does more than colour perception. It can select an action that changes which data become available and how the result is explained. The forecast reaches into the experiment.

Surprise arrives with paperwork

In the cleanest story, an expectation is contradicted and the system updates. You expect catastrophe in an elevator. You ride to the fourth floor. Nothing catastrophic happens. Prediction error: the outcome was safer than predicted.

But prediction error is an opportunity for learning, not a command. Human models are protected by explanations. “It was only four floors.” “Someone I trust was with me.” “I had medication in my pocket.” “I stood beside the door.” “I got out before the feeling peaked.” Each statement may be factually true. Together they can prevent the safe outcome from testing the broad proposition that the elevator is intolerable.

This is not evidence of stupidity. Every stable model needs some resistance to anomalies. If one warm January afternoon erased the concept of winter, learning would be chaos. Scientists do not abandon a mature theory because a single instrument produces an odd reading; they inspect the instrument, the conditions, and rival explanations. Everyday cognition needs the same steadiness.

The problem begins when a threat model becomes an unbeatable diplomat. The friendly dog is an exception. The kind audience was unusually kind. The normal scan was too early. The safe flight was lucky. Contradictions are welcomed at the border, photographed, and denied citizenship.

What determines whether surprise changes the model? At least four things matter. The outcome must be noticed. It must be relevant to the actual prediction. It must be credible rather than dismissed as noise. And its cause must be attributed in a way that bears on the model.

Repeating an experience without examining these conditions may produce endurance without much revision.

The converse matters too. One frightening outcome can be over-credited. If panic occurs on a train, the system may learn “trains are dangerous,” although the panic was influenced by sleep loss, caffeine, a recent illness, and an already elevated level of stress. The event is real; the causal lesson is not self-interpreting.

When uncertainty itself receives the microphone

Sometimes the high-precision message is not “danger is present.” It is “not knowing is dangerous.” This distinction explains why a person may continue to check after the most likely threat has been answered.

Did I lock the door? Probably. Can I be certain? No. Is my memory trustworthy? Usually. Can I be certain that *this* memory is trustworthy? Again no. The target moves upward from the state of the door to confidence in memory, then to confidence in confidence. First-order evidence cannot finish a second-order inquiry whose stopping rule is perfect certainty.

This form of doubt is not confined to checking. A medical result can answer one diagnostic question and leave open whether the test could have missed something. A loving reply can settle whether a partner is angry now and leave open whether the relationship will last. An excellent performance can show that one presentation went well and leave open whether the next one will.

The future is structurally unable to provide the guarantee demanded of it. When uncertainty itself has acquired high threat value, the absence of closure becomes evidence for continued control. The system need not believe that the worst outcome is probable. It need only treat unresolved possibility as sufficient reason to act.

Large studies do find a strong association between intolerance of uncertainty and worry, but the relation is not a pure window onto one hidden mechanism. It weakens when overlapping depression and questionnaire content are accounted for, and it differs across presentations. ^[UW-01] The empirical point is not that every anxious

person “cannot tolerate uncertainty.” It is that, for some people in some states, uncertainty is not merely an informational gap. It functions as a cost.

Two confidence systems in one person

“I know this is probably safe,” someone says, while standing ten metres from the door they cannot enter. The sentence is often interpreted as proof that the person is irrational or unmotivated. It may instead reveal a disagreement among levels of control.

Propositional confidence lives in language: what I endorse when asked. Sensorimotor confidence appears in preparation, attention, muscle tone, and the actions that remain reachable. A person can sincerely report a low probability of harm while the action-selection system places a high penalty on approach. There is no need to imagine two tiny people arguing in the head. Different learning histories, time scales, bodily signals, and contexts can exert different weights on one decision.

This is also why explanation can be liberating without being sufficient. Learning the physiology of panic may remove shame and reduce mystery. Yet knowledge acquired in a calm room is not automatically granted decisive weight in a crowded train. The issue is not whether the information entered memory. It is whether it can participate in control under the conditions in which the old prediction becomes loud.

Precision is thus situational. Sleep loss, pain, social isolation, unfamiliarity, intoxication, withdrawal, and acute stress can alter what the system trusts. A coping strategy that is obvious on Sunday afternoon may be functionally absent at 3:00 a.m. This does not make the Sunday insight false. It shows that access to evidence is part of the phenomenon to be explained.

Can the system still be surprised by safety?

A healthy threat expectation need not be weak. It needs terms under which it can change. That is the practical test hidden inside the technical language.

What exactly is predicted? How far does the prediction extend? Which observations would count against it? How many exceptions can be explained away? What role do protective actions play in the explanation? Does the person update when calm but not when aroused? Can a safe outcome become more than a lucky anecdote?

These questions move us beyond the demand to “think positively.” They also prevent predictive language from becoming decorative neuroscience. A prior is not a curse. Precision is not a diagnosis. Prediction error is not a guaranteed cure. They are tools for asking why certain futures acquire authority and why certain present facts fail to dislodge them.

At the edge of this chapter sits a further complication. Evidence never arrives attached to only one possible cause. The palpitation could reflect danger, exertion, caffeine, fever, medication, a benign rhythm variation, or panic. The delayed message could mean rejection, concentration, poor signal, sleep, or nothing of lasting significance. Updating requires not only a prior and a new observation, but a contest among explanations.

That contest is where an eighteenth-century minister named Thomas Bayes unexpectedly enters the story.

Chapter 18 - Bayes and the Committee of Possible Causes

An equation with no bedside manner

Thomas Bayes died in 1761, long before search engines, cardiac wearables, and the blue ticks that indicate a message has been read. It seems unlikely that he expected his name to appear in a book about panic. Yet the logic attached to it is unusually helpful precisely because anxious evidence so often arrives without a label.

Bayes' theorem is commonly written:

$$P(H | E) = \frac{P(E | H)P(H)}{P(E)}$$

Here, H is a hypothesis and E is evidence. The expression says that our revised confidence in a hypothesis after seeing evidence depends on where we started and on how expected that evidence would be under competing explanations. The symbols are compact. The world they describe is not.

Suppose the hypothesis is “this tightness signals a dangerous cardiac event” and the evidence is “my chest feels tight.” A dangerous cardiac event can cause tightness. That fact matters, but it is not the whole calculation. How common was the event before the sensation occurred? How often does tightness arise from other causes? What other symptoms, risks, tests, and circumstances are present? Evidence does not become diagnostic merely because it is compatible with a frightening hypothesis.

Nor should this example be used to reassure a stranger that chest symptoms are “only anxiety.” New, severe, or concerning physical symptoms deserve appropriate medical assessment. Panic and medical illness can coexist, and panic disorder has been associated at the population level with later coronary disease; an association does not tell us the cause of one person's pain. ^[DD-05] Bayesian humility cuts

both ways. It resists catastrophe without turning uncertainty into dismissal.

The missing denominator

Online symptom searches often reveal statements of the form “Condition X can cause symptom Y.” That is $P(E|H)$: how compatible the evidence is with the hypothesis. What the sentence usually does not tell you is how common the condition was to begin with, how often the symptom appears without it, and which alternative explanations fit at least as well.

The omitted denominator is where many alarming stories hide their ordinariness. A rare disease can produce a common symptom. Most people with the common symptom may still not have the rare disease. This is the base-rate problem, and base rates suffer from a public-relations disadvantage. They are usually written as pale numbers. Case stories have names, photographs, and plot.

A vivid exception can therefore acquire more psychological force than a large background statistic. That bias is not unique to anxiety. It is part of ordinary human judgment. Anxiety can make the bias behaviourally expensive: another search, another scan, another hour spent comparing images, another promise that this will be the final check.

The inverse error also exists. A common benign explanation does not make a serious alternative impossible. Good reasoning is not the ritual repetition of “probably nothing.” It is proportional updating: weighing prevalence, pattern, context, and consequences without making either possibility or probability carry more than it can.

Bayesian language earns its place here by forcing a small but important separation. “Could cause” is not “probably caused.” “Fits” is not “distinguishes.” “Rare” is not “impossible.” “Common” is not “certain.” Those distinctions are less emotionally satisfying than a verdict, but they are the grammar of honest inference.

Competing stories, not evidence versus ignorance

Consider a manager who writes, “Could we talk tomorrow?” One model says a problem has arisen. Another says the manager wants an update. Another says schedules are being changed. Another says the message was written quickly between meetings and contains almost no useful information at all.

Anxiety is sometimes described as believing bad things without evidence. More often, evidence exists but is radically underdetermining. The message really was brief. The manager really did omit a smile or explanation. Those observations support several stories. The struggle concerns which causal model is available, vivid, and trusted.

The hypothesis space matters. If only two explanations are considered-“everything is fine” or “I am being fired”-the threatening story may dominate because false reassurance feels naive. Add mundane intermediate possibilities, and the same evidence occupies a different field. This is not positive thinking. It is causal pluralism.

Yet generating alternatives can itself become a ritual. Someone may write forty possible explanations and demand a probability for each. The exercise then serves certainty rather than flexibility. The goal is not to build an encyclopaedia of harmless causes. It is to prevent one underdetermined observation from appointing one cause as sovereign before the world has supplied enough information.

This point becomes particularly important after safe experiences. A flight lands without incident. Did this occur because flying was safer than predicted? Because a ritual prevented disaster? Because the person remained hypervigilant? Because a companion was present? Because this flight was the exception? One outcome can be recruited by several models. Contact with the situation is therefore not automatically corrective. Causal attribution decides what the contact teaches.

Probability is only half a decision

Imagine two people who both estimate a one-percent chance of a bad outcome. One proceeds; the other withdraws. Neither need have miscalculated probability. Their decisions may differ because they assign different costs to the event, different value to the opportunity, and different confidence to their ability to cope.

This is where inference opens into decision theory. A crude piece of bookkeeping might compare the expected losses and gains of several actions, together with effort, uncertainty, and opportunity cost. No one claims that the brain solves a neat spreadsheet before entering a lift. The bookkeeping is useful because it exposes variables concealed by the sentence "I know it probably won't happen, but I still can't do it."

Perhaps the person thinks the feared outcome would be unbearable. Perhaps a small chance of harming someone carries immense moral weight. Perhaps the action offers little reward. Perhaps leaving is cheap. Perhaps arousal has temporarily made long-term value difficult to feel. The probability estimate may be quite reasonable while the decision remains defensive.

Many anxious simulations are detailed until the moment of disaster and strangely blank afterward. Rejected-then nothing. Panic on the train-then nothing. Make a mistake at work-then the screen goes dark. The event has a rich model; recovery does not.

But lives continue after peaks. There may be embarrassment, repair, medical care, a phone call, time, learning, apology, adaptation, grief, help, or a changed plan. None of this guarantees a happy ending. It restores a missing variable: the organism's capacity to act after the unwanted event. Underestimating coping can inflate expected loss without changing the estimated probability of danger at all.

The internal prosecutor

Human costs are not confined to injury and money. They include shame, responsibility, identity, and anticipated self-judgment.

“If something happens, I will never forgive myself.” Now the future contains not only an adverse event but a trial conducted by the future self. The person may avoid a low-probability risk because the anticipated verdict—careless, selfish, incompetent—feels intolerable.

This is especially visible in decisions involving loved ones. A parent can know that no childhood is risk-free and still experience every preventable possibility as a moral summons. Worry then feels like evidence of love; vigilance feels like proof of responsibility. To relax is not merely to lower an estimate. It is to risk becoming the sort of person who failed to care enough.

Moral seriousness should not be mocked out of the model. Some decisions do involve responsibility, and regret can be devastating. The mistake is to assume that unlimited anticipatory suffering purchases unlimited protection. Often it does not. It may instead obscure the distinction between a useful precaution and a ritual designed to silence the internal prosecutor.

Decision theory helps by making the hidden tribunal visible. What outcome is being prevented? What probability is assigned to it? What would make it costly? How much of that cost is predicted self-condemnation? What does the defensive act protect, and what else does it take away? These are different questions. Anxiety compresses them into one word: danger.

Ambiguity refuses to sit still for arithmetic

Many consequential choices do not supply known probabilities. Will this relationship last? Will the treatment help? Will a new job be worth the move? Will the political situation worsen? We do not possess a casino's fixed odds. We are uncertain about the probabilities themselves, and sometimes about the list of possible outcomes.

This distinction is often called ambiguity. Under ambiguity, more calculation does not necessarily produce more knowledge. It may produce decimals attached to guesses. Numbers can clarify assumptions, but they cannot manufacture a stable frequency where none exists.

Someone who experiences ambiguity as highly costly may prefer a familiar bad option to an uncertain promising one. The known job is painful, but its humiliations have a schedule. The new role could be better, but its social world is uncharted. A predictable small life can defeat a possibly larger one because predictability is part of the value calculation.

At this point, Bayesian vocabulary can become a costume worn by ordinary fear. One can assign invented probabilities to every branch and call the result rational. The more honest use of the framework is diagnostic rather than oracular: to ask which assumptions control the answer, how sensitive the conclusion is to those assumptions, and which observation would actually discriminate among models.

Some uncertainty is reducible. Ask when the train leaves. Check whether the medication interacts with another one through a qualified source. Obtain an appropriate medical evaluation. Some uncertainty is irreducible at the moment of choice. No amount of checking can establish how a first date will feel before it occurs. Wisdom includes knowing which kind of uncertainty one is trying to solve.

Active organisms alter the evidence

Bayesian diagrams are often drawn as though the agent sits still while evidence rains down. Living creatures do not wait so politely. They move their eyes, seek people, open messages, avoid streets, ask questions, leave rooms, and build environments in which some observations become likely and others nearly impossible.

Formal active-inference theories attempt to describe the joint problem of perception and action. They are mathematically rich and scientifically contested in their scope; clinical metaphors should not outrun the evidence. The modest principle we need is simpler: a policy changes the data that will later be available to update the policy.

Avoid every dog and you collect little fine-grained evidence about dogs. Rehearse every conversation and successful conversations may be credited to rehearsal. Check a bodily metric every few minutes and ordinary variability becomes a dense stream of potential anomalies.

Ask a loved one for certainty and receive evidence about what they say under questioning, not necessarily about what the relationship feels like without repeated testing.

Action also directs attention. Enter a room intending to detect rejection and the room supplies eyebrows, pauses, and ambiguous glances. This does not mean the observations are imaginary. It means sampling is selective. A model can make itself look accurate by choosing a narrow channel through which the world is allowed to answer.

Avoidance is the most decisive form of sampling because it can remove the test entirely. “I cannot travel alone” may remain true in the person's practical model not because solitary travel failed, but because no observation was permitted. Absence of disconfirmation is then mistaken for confirmation.

What would make the model lose?

Good inference requires more than collecting supporting evidence. It requires a possible defeat.

Suppose the model says, “If I do not check the lock, the house may be burgled.” A night without burglary cannot prove that checking was unnecessary; burglaries are rare. But if the model also says checking is the only thing preventing disaster, then the checking itself makes the causal claim difficult to test. A serious experiment must identify what would count as evidence for the competing explanation and what level of risk is ethically acceptable.

The same standard applies to theory. Predictive and Bayesian words can explain almost anything after the fact if priors, precision, costs, and policies are adjusted freely. A framework that cannot be wrong is not deep; it is elastic. Its value lies in generating discriminating questions: Does a person overestimate probability, underestimate coping, overprice ambiguity, assign safety to a ritual, or fail to retrieve alternatives under arousal? Those possibilities predict different patterns.

Bayes does not tell a frightened person what to do. It tells us why one sensation, message, or safe flight cannot interpret itself. It asks us to keep prior plausibility, evidential fit, causal attribution, cost, and action separate long enough to see which one is governing the decision.

Then action produces a consequence. And consequences have their own way of teaching.

Chapter 19 - Reinforcement Learning Under Threat

The reward nobody puts on a poster

A laboratory rat presses a lever and receives food. The scene is so familiar that the animal has become an unpaid ambassador for reinforcement learning. Humans are more difficult. We ask whether the lever is meaningful, whether pressing it reveals a character flaw, and whether another lever has better reviews.

Still, the central fact survives the comedy: actions have consequences, and consequences alter the probability of future actions.

Anxiety has a particularly effective consequence at its disposal. Relief.

Decline the invitation; the stomach loosens. Leave the queue; the heart slows. Check the result; uncertainty drops. Ask once more whether everything is all right; warmth returns to the room. The action removes or reduces an aversive state. In behavioural language, it can be *negatively reinforced*. “Negative” means subtraction, not badness: something unpleasant is taken away, and the action becomes more valuable.

Relief is not fake. It is not a foolish prize awarded by a defective brain. It is a genuine state change, often rapid and reliable. That is precisely why avoidance can be so persistent. Advice that lists the long-term disadvantages while ignoring the immediate payment misunderstands the market in which the defensive action is competing.

One small equation, unpacked

A minimal reinforcement-learning rule updates the expected value of an action after its outcome:

$$Q_{t+1}(s, a) = Q_t(s, a) + \alpha \delta_t$$

and a common form of the prediction error is:

$$\delta_t = r_t + \gamma \max_{a'} Q_t(s', a') - Q_t(s, a).$$

The symbols are less forbidding once they are given jobs. $Q_t(s,a)$ is the value currently assigned to action a in state s : for example, the learned value of leaving a crowded shop while highly aroused. The outcome r_t is the immediate result, which may include a sharp fall in distress. The term involving γ represents expected future value; γ determines how strongly later outcomes matter. The prediction error δ_t is the difference between what the action delivered and what the system expected. Finally, α , the learning rate, determines how much this surprise changes the old value.

This is a model, not a claim that a single neuron writes equations in chalk. Its usefulness lies in the questions it makes unavoidable. What counts as the reward? Which state is the learner in? How far into the future is value carried? What was expected? Which action receives credit?

Suppose leaving the shop produces more relief than expected. The prediction error is positive relative to the goal of reducing distress, so the value of leaving in that state rises. Nothing pleasurable needs to occur. A bad afternoon can teach a strong lesson if one action makes it less bad quickly.

Now notice what the equation does not automatically include. It does not know that repeated escape may shrink the person's world in six months unless those delayed costs are represented and linked to today's action. It does not know what would have happened had the person stayed. The unchosen action supplies no direct outcome. Reinforcement learning is powerful partly because it learns from what happened; anxiety can persist partly because the most important evidence concerns what was never allowed to happen.

Excellent cash flow, disastrous accounts

Avoidance resembles an investment that pays a small dividend at once while quietly consuming the capital.

Cancel the party: relief in thirty seconds. Attend: anxiety first, perhaps an awkward opening, perhaps pleasure later, perhaps a friendship whose value cannot be known for months. Check the email again: a second of certainty now. Resist checking: no immediate celebration, only a gradual possibility that uncertainty becomes more bearable. Avoid exercise: bodily alarm falls today. Continue exercising appropriately: fitness, confidence, and new interoceptive learning may emerge slowly.

The defensive policy has excellent liquidity. Its benefits are close, vivid, and causally obvious. Long-term losses are delayed, distributed, and hard to attribute. No single cancelled gathering creates loneliness. No single avoided road eliminates driving skill. No single reassurance request dissolves trust in one's own judgment. The debt arrives as a changed life rather than an itemized invoice.

This temporal asymmetry helps explain how locally sensible decisions can create a globally painful trajectory. On each occasion, leaving may be the easiest way to reduce a state that feels urgent. Over time, the person has fewer experiences of remaining, adapting, recovering, or finding unexpected value. The action landscape changes, and then the world really does become harder to enter.

It would be tempting to call this “choosing short-term reward over long-term reward” and move on. But the phrase can sound moralistic, as though the person lacks patience. Under threat, the later reward may not merely be discounted; it may be inaccessible to imagination. Connection next month cannot compete with suffocation now if the threatened system cannot represent next month with enough force to influence action.

Credit assignment: who saved the day?

Consequences teach only after someone-or some learning process-decides what caused them.

A person boards a plane after touching the door three times, carrying a particular object, and monitoring every sound. The plane lands safely. Which action receives credit? Aviation engineering is one candidate. The ordinary safety of the flight is another. The ritual is a third. If relief follows the ritual before take-off and safety follows the whole sequence, the ritual can acquire value without ever having controlled the aircraft.

Laboratory evidence shows that avoidance or safety responses can preserve threat expectations under some conditions, though not every safety behaviour is harmful and the context matters. ^[LE-06] The lesson is not “remove every support.” It is to ask whether support makes contact possible or whether it monopolizes the explanation for survival.

Credit assignment also affects social reassurance. If another person repeatedly confirms that a feared interpretation is wrong, the outcome may teach “I can trust this person to settle uncertainty,” not “I can let an uncertain thought remain unresolved.” The relationship provides real regulation. It may widen action or become the only authorized source of safety. Similar behaviour can serve opposite learning functions.

The causal story is often invisible. A person says, “I managed the event,” while privately adding, “because I sat near the door.” The event looks like exposure from the outside. Inside the learning system, the door may have won the medal.

Pavlov's vote in an instrumental election

Reinforcement learning is commonly described as instrumental: choose an action, observe an outcome, update its value. But action selection is not governed by one controller. Older Pavlovian tendencies cast votes too.

A cue predicting reward can invigorate approach. A cue predicting punishment can favour inhibition, withdrawal, freezing, or escape. These biases are often adaptive. When a predator is close, a lengthy debate about personal values is poor design. Yet a Pavlovian response

can conflict with the instrumental action that would produce the better outcome.

Open the medical letter to obtain useful information; the threat cue says freeze. Approach the podium to complete a valued talk; social alarm says withdraw. Remain with a rising bodily sensation long enough to learn its course; panic says escape. Speak honestly in a relationship; attachment threat says appease or disappear.

In experimental tasks, researchers sometimes examine *Pavlovian congruence*: whether action is easier when the Pavlovian bias and the instrumental requirement point in the same direction. “Act to gain reward” and “withhold action to avoid punishment” are often congruent. “Act to avoid punishment” can be harder because threat promotes inhibition even though action is required. Human findings across anxiety and mood symptoms are mixed, dependent on task and sample. There is no single Pavlovian defect that identifies an anxious person.

The conceptual gain remains. Failure to approach is not necessarily a cold calculation that approach lacks value. It may be the outcome of competition among controllers operating at different speeds. The action that a person endorses in principle can lose at the moment a fast defensive bias changes the field.

A state is part of the action

The notation $Q(s,a)$ includes a detail ordinary advice often omits: value belongs to an action *in a state*.

Calling a friend while mildly uneasy is not the same learned problem as calling while trembling after three nights of poor sleep. Entering a shop at opening time is not identical to entering at rush hour. The system may have learned several values for what language calls the same action.

This is why a strategy can be available in the therapist's office and vanish at the station. “I know what to do” may be completely true in one state. Under intense arousal, working memory, attention, timing,

bodily interpretation, and expected relief shift. The effective menu contracts.

Questionnaires administered in calm conditions can therefore overestimate behavioural flexibility under perturbation. Asking what someone *would* do is not identical to observing which policy remains reachable when threat is high. A serious account of recovery must examine state transitions, not only declared preferences.

The same point protects against simplistic celebration when a person succeeds once. One easy-context success is valuable but may not yet generalize. Learning must travel across time, bodily states, places, companions, and levels of uncertainty. Otherwise the alternative remains a specialist, licensed to work only on quiet afternoons.

There is no universal anxious learning rate

Computational psychiatry offers an irresistible fantasy: perhaps anxiety is one parameter. Learn too quickly from punishment, too slowly from safety, explore too little, and there it is—a neat numerical fingerprint.

The literature has not granted that wish. Results differ across tasks, diagnoses, symptom dimensions, models, and states. A parameter called “learning rate” in one experiment may not measure the same psychological process as a similarly named parameter in another. Group averages can also conceal distinct subgroups or trajectories, as fear-learning research has repeatedly warned. ^[LE-08]

Heterogeneity is not a disappointing footnote to be hidden until the end. It changes the theory. Some people may update rapidly after negative outcomes. Others may learn normally but assign excessive cost to ambiguity. Some may acquire threat readily; others mainly fail to retrieve safety in a new context. Some may seek information intensely while avoiding consequential action. Anxiety is unlikely to be a single broken dial.

This also means that a plausible equation is not evidence that its parameters have been measured in a particular person. Computational

models simplify behaviour so that competing explanations can be compared. They are instruments for disciplined uncertainty, not high-tech personality tests.

Alternatives must be competitive at the moment of sale

If avoidance pays in seconds, an alternative cannot compete merely by being admirable in theory.

“Stay and breathe” may exist as a sentence but not as a reachable policy at peak arousal. “Allow the uncertainty” may be remembered only after reassurance has already been requested. “Act according to values” may point in the right direction while lacking the practice, timing, environmental support, and immediate credibility needed to defeat escape.

For an alternative to become competitive, it must be sampled in relevant states and receive interpretable outcomes. It may need to be smaller than the heroic version imagined by advice. Remaining for ninety seconds can be a policy. Opening the message and waiting before replying can be a policy. Entering with support that facilitates contact, then gradually changing the role of that support, can be a policy.

The aim is not to make distress disappear instantly; if disappearance remains the only reward, the old route retains home-field advantage. Alternative action can be reinforced by other outcomes: completing what mattered, discovering that arousal changes, recovering after awkwardness, learning which risks are real, preserving self-respect, or enlarging tomorrow's available options.

Practice also changes latency. A response rehearsed only in calm reflection may arrive too late. Repeated, appropriately graded use can make it faster and less cognitively expensive. In learning terms, recovery is not only acquiring a good idea. It is changing the contest so that the good idea can enter the race before the defensive policy has crossed the finish line.

Flow Hijacked synthesis: when relief captures the route

The following proposal is a **Flow Hijacked synthesis**, not an established reinforcement-learning finding.

Anxious capture may occur when three processes cooperate: rapid relief gives a defensive action high effective value; that action prevents or distorts the observations that could revise the threat model; and arousal reduces access to slower, value-rich alternatives. The loop can then sustain itself without any single abnormal learning parameter. Relief strengthens the route, restricted sampling protects the forecast, and the protected forecast makes relief valuable again.

This proposal should be testable. It predicts that measuring symptom intensity alone will miss important variation. We would need within-person observations of state, available actions, avoidance cost, relief timing, causal attribution, and later opportunity. It also predicts that reducing distress without widening action may leave the capture intact, while a modest symptom change accompanied by greater behavioural reach could represent meaningful change.

Evidence on costly avoidance supports the possibility that some people with anxiety disorders continue avoiding even when approach carries competing rewards, but the pattern is heterogeneous and the available studies do not validate the Flow Hijacked construct. ^[LE-07] If avoidance does not predict later contraction after symptoms and real-world constraints are considered, or if contraction reliably precedes avoidance, this part of the synthesis would need revision.

The model is useful only if it can lose. Otherwise “capture” is a dramatic name for whatever happened.

The unchosen action has no testimony

Reinforcement learning explains how consequences teach. It also reveals a silence at the centre of anxiety: only selected actions generate direct outcomes.

Leave the party and you learn that leaving brings relief. You do not learn how anxiety would have changed after twenty minutes, whom

you might have met, or whether an awkward moment could have been repaired. Check the symptom and you learn that checking briefly alters uncertainty. You do not observe the trajectory of the urge when left unanswered.

The missing outcome is counterfactual. It belongs to the road not taken. A model can estimate it, but estimation built from years of avoidance becomes increasingly self-referential. “Staying would have been terrible” remains plausible because staying never receives a witness.

To learn about an unchosen action, the system must sometimes choose it. That is not a demand for recklessness. It is the unavoidable price of evidence. The next chapter is about how organisms decide when the price is worth paying-and what happens when a life becomes very good at exploiting the same protective answer.

Chapter 20 - The Price of Not Exploring

At the edge of the playground

A child stands beside a parent at the edge of a playground. Staying close offers known safety. The climbing frame offers uncertain footing, loud children, the possibility of embarrassment, and a respectable chance of a scraped knee. It also contains nearly everything the child cannot obtain by staying at the edge: balance, friendship, surprise, competence, play.

This is the explore-exploit dilemma before anyone gives it a Greek letter. Exploit what is already known to work, or spend resources sampling an option whose value is uncertain. Neither policy is always superior. When danger is high, resources are scarce, or the time horizon is short, relying on a known option can be wise. When the horizon is long, endless exploitation can trap an agent inside a solution that is safe enough and far smaller than what is possible.

Exploration charges an uncertainty fee. The organism pays effort and accepts that some outcomes will disappoint. In exchange it acquires information that prediction alone cannot supply.

Anxiety can make the fee feel unaffordable. Immediate safety rises in value, future reward becomes less vivid, and the familiar action arrives with a receipt proving that it worked yesterday. The unknown option has no such paperwork. Its proof comes after the action.

Menus are not meals

Imagine reading the menu of a new restaurant. You can inspect ingredients, reviews, photographs, hygiene scores, parking, and the biography of the chef. At some point the information curve flattens. One fact remains unavailable: what this evening, with these people and this version of you, will actually be like.

Exploration produces observations that could not be fully specified in advance. Even a mediocre meal supplies something the menu could

not. The evening may reveal a preference, a tolerance, a conversational rhythm, a route home, or simply that disappointment is survivable.

An anxious decision system may count the mediocre meal as evidence that exploration was a mistake. But if only excellent outcomes justify sampling, learning becomes impossible. Experiments are valuable partly because their results are unknown. A life designed to prevent every unprofitable trial also prevents discovery.

This does not mean every experience is secretly good. Some restaurants are awful; some dates are cruel; some jobs are exploitative. Exploration is not optimism with a backpack. It is a controlled willingness to acquire information while retaining the right to update, leave, protect, and choose again.

The researcher who never runs the experiment

Anxiety does not simply reduce exploration. It can produce prodigious exploration in one space and almost none in another.

Read fifty hotel reviews. Compare twelve routes. Ask five friends. Watch street videos. Examine crime maps. Open twenty-seven tabs. Cancel the trip.

The information search is real. It may be careful, intelligent, and exhausting. Yet its function can differ from behavioural exploration. One kind of inquiry seeks enough information to make a reasonable choice. Another seeks to abolish the uncertainty that only action can resolve.

This distinction helps make sense of mixed experimental findings. Cognitive worry can sometimes accompany directed information seeking; acute or bodily anxiety can suppress sampling in other contexts. Asking whether “anxiety increases exploration” is too coarse. We need to ask what is being explored, in which state, at what cost, and toward what end.

A person may be adventurous among explanations and conservative among lives. The hypothesis space expands while the action space contracts. Search becomes movement without travel.

The confidence certificate paradox

People often wait for confidence as though it were an entry document issued before experience.

I will speak when I feel like a speaker. I will travel when I know I can handle panic. I will date when rejection no longer terrifies me. I will make the decision when I am certain it is correct.

But much of confidence is retrospective. Speakers become experienced by speaking while inexperienced. Travellers learn coping by encountering the unplanned. Decision confidence grows partly through making imperfect choices, repairing errors, and discovering that uncertainty did not revoke personhood.

The prerequisite is often a consequence.

This does not require the motivational slogan “feel the fear and do it anyway.” Sometimes fear carries information, and sometimes the proposed act is too large, too fast, or genuinely unsafe. The practical question is smaller: what experiment could produce the missing evidence without demanding certainty first?

Competence and confidence are also distinct. Someone may possess the skill to drive and lack confidence after a frightening event. Another may feel confident and lack skill. Exploration should be designed to learn which problem exists, not to use courage as a substitute for training.

Exploration updates the self-model

When you enter an uncertain situation, you learn about the situation. You also learn what kind of organism you are within it.

Can I improvise when the plan breaks? Can I feel embarrassed and remain? Can I ask for help without collapsing into shame? Can my

heart become loud and then quiet? Can I leave deliberately rather than flee? Can I care for myself after a bad outcome?

These are empirical questions. Introspection can propose answers, but it cannot settle them alone. A self-model built only from imagined trials will tend to recycle old evidence. New action gives the person an opportunity to discover capacities that were not available, needed, or visible in the original learning environment.

The discovery may be modest. “I enjoyed it” is not required. “I stayed long enough to choose rather than bolt” may be new information. “It was difficult, and I recovered by evening” restores time to a simulation that previously ended at the peak. “I asked for accommodation and was treated decently” can revise both the social model and the model of agency.

Exploration can also reveal genuine limits. The route is inaccessible. The workplace is harmful. The symptom needs medical attention. A person is not obligated to transform every boundary into a challenge. Accurate self-knowledge includes constraints as well as capabilities.

Avoidance can manufacture part of the forecast

If driving is avoided for years, driving may become objectively harder. If exertion is avoided, fitness can decline. If social contact disappears, fluency may weaken. If every decision is outsourced, independent choice becomes less practiced.

Now the system encounters real evidence of reduced capacity. “See? I cannot do it.” The sentence is no longer wholly forecast. Behavioural restriction has altered the organism and environment in a direction that partially confirms the original fear.

This is a painful feedback loop because it can be mistaken for proof that avoidance was wise from the beginning. In fact, present difficulty may be jointly produced by the original threat, the protective response, lost practice, and changed opportunity. Causal history matters.

The loop is not inevitable. Skills can be rebuilt, and not every avoided capacity decays. Nor should real disability be redescribed as anxiety-created fragility. The point is conditional: some protective actions change the future costs they were intended to prevent.

Costly-avoidance experiments illustrate that competing rewards do not automatically overcome learned protection, and that people within the same diagnostic group can show quite different patterns. ^[LE-07] A laboratory task cannot tell us why one person's world has narrowed. It does warn against assuming that making the reward obvious is enough.

Horizons decide what counts as rational

Explore-exploit decisions depend on time. If a café is closing in five minutes, sampling every pastry is absurd. If you will eat there daily for ten years, trying something new may be worth one disappointing breakfast.

Threat can shorten the effective horizon. The next sixty seconds dominate the next six months. A choice that looks irrational from a distance can be rational inside the narrow temporal window created by panic. Conversely, advice that invokes distant flourishing may fail because flourishing has too little present weight.

Widening the horizon is not simply telling someone to think long term. The future must become affectively and behaviourally reachable. Its rewards need enough specificity to compete: not “have a better life,” but “be able to take my daughter to the sea,” “speak at the meeting I care about,” or “sleep without employing the phone as a night watchman.”

Even then, delayed value should not be used to bully the present self. Sustainable exploration respects the nervous system's current capacity. A sequence of tolerable experiments can generate more learning than one heroic exposure followed by months of retreat.

When the road is blocked, not feared

Any theory that praises exploration risks becoming cruel. Poverty restricts choices. Disability changes the cost and accessibility of environments. Racism, antisemitism, sexism, homophobia, transphobia, and other forms of hostility make some social forecasts more accurate than comfortable observers wish to admit. Caregiving can make spontaneity impossible. Immigration status, war, illness, transport, work schedules, and unsafe housing can close real roads.

Calling every restriction “avoidance” transfers responsibility from the world to the person. It can convert structural harm into a homework assignment.

The relevant distinction is not between brave and fearful people. It is between paths closed primarily by external conditions, paths closed primarily by protective prediction, and the many paths closed by both. External support, accommodation, medical care, money, policy, companionship, and safety can be prerequisites for exploration rather than signs that it is inauthentic.

Support should be judged by function. Does it allow a person to encounter meaningful uncertainty and gather new evidence? Or does it require the world to remain perfectly controlled before any action can occur? The same companion, medication plan, route, or preparation can operate differently at different stages. Context defeats slogans.

A map with a thousand roads and three footprints

Prediction extrapolates from what is known. Exploration changes what can be known.

Without sampling, a model can become exquisitely accurate about a tiny territory. It knows every exit, every reassurance source, every bodily fluctuation, every reason the unfamiliar route may fail. Efficiency rises while reach falls.

The map may contain a thousand roads. Three have footprints.

This is how a future can become small without anyone choosing a small life. Each local decision can be understandable. The contraction appears only when the decisions are viewed across time.

The **Flow Hijacked synthesis** proposed later will treat this loss of reachable alternatives as more than a metaphor, but it must not be smuggled in as established fact. Here the narrower claim is enough: exploration supplies information, skill, and self-knowledge that exploitation cannot; defensive learning can make that information costly to obtain; and real conditions determine which experiments are possible and humane.

We now leave the machinery for a while and look at the lives it can shape. There is no single anxious future. Worry, panic, social fear, compulsive doubt, phobic restriction, and trauma-related vigilance overlap, but they are not interchangeable costumes on one disorder. A good synthesis must survive their differences.

PART V

MANY FACES OF AN ANXIOUS FUTURE

Shared processes without erasing different lives and diagnoses.

Chapter 21 - The Future With No Closing Time

Worry changes the subject

Generalized anxiety has an unusual talent: it survives solved problems.

The medical result is reassuring. For twenty minutes the room becomes quiet. Then money takes the chair. The invoice is paid; a child's silence acquires significance. The child calls; work has been waiting patiently in the corridor. Work settles; now there is sleep, and if sleep arrives too easily there may be time to wonder why one has become so careless.

A specific fear can be fiercely loyal to its object. Chronic worry has wider tastes. Its continuity does not depend on any one catastrophe remaining plausible. Content changes while an operation persists: search the future for what has not yet been secured.

This is easy to misunderstand because many worries are reasonable in isolation. Bills matter. Children can need help. Bodies become ill. Work contains consequences. Generalized worry does not require bizarre subjects. It can be built from ordinary responsibilities processed without a reliable stopping rule.

Planning has stopping rules. A bill is due Friday; the transfer is scheduled; there is enough in the account. More thought may not be pleasant, but the task has a criterion for completion. Worry often uses another criterion: stop when nothing bad remains possible. The future cannot meet it.

An open system is placed under a closed-world demand. No wonder the office stays open all night.

The factory and the product line

A single-alarm metaphor fits a single looming danger. Generalized worry is better imagined as a factory whose product lines can change without interrupting production.

What if? Then what? What would that mean? What might I have missed? How would I forgive myself for not thinking of it?

One question generates another because each answer contains an edge. If tomorrow's presentation is prepared, perhaps sleep will fail. If sleep is managed, perhaps the preparation will sound rehearsed. If the delivery succeeds, perhaps expectations will rise. Worry solves a local uncertainty and uses the solution to discover a new one.

The factory is not powered only by fear. It may be powered by conscientiousness, imagination, love, and intelligence. People who worry chronically are sometimes extremely capable at identifying dependencies and planning contingencies. The pain lies in a useful faculty being recruited past the point of useful work.

This helps explain why generic advice to “stop overthinking” can feel both obvious and insulting. Thinking is not a tap turned on by mistake. It is doing a job. The therapeutic and scientific question is what job the process believes it must continue doing after practical planning has ended.

When worry dresses as responsibility

Parents worry because they care. Clinicians rehearse rare complications because people depend on them. Workers imagine mistakes because standards matter. To stop worrying can therefore feel less like rest than negligence.

The hidden rule may be: if I suffer in advance, I have taken the danger seriously. Worry becomes a receipt for responsibility.

But anticipatory suffering and preventive action are not the same currency. Five more hours of mental rehearsal may not improve the plan. It may impair sleep, attention, patience, and the very performance the worry is meant to protect. The fact that worry is costly does not prove it is effective. Cost can be mistaken for commitment.

There is also magical responsibility in a less theatrical form. “The one time I relax will be the time something happens.” No one

necessarily believes thoughts control events. Rather, relaxation removes the feeling of preparedness, and the contrast is experienced as exposure. Bracing appears prudent because it prevents the shock of being caught emotionally unarmed.

The humane response is not to persuade people to care less. It is to separate care from continuous internal punishment. Responsibility asks what can be done. Worry often asks how long one must remain distressed to prove that the doing was serious.

Paying the emotional bill in advance

The Contrast Avoidance Model proposes that worry may sometimes help a person maintain a negative emotional state and thereby reduce the jolt of a sudden shift from calm to distress. If you never fully unbrace, bad news cannot knock you down from as great a height.

Experimental and prospective work provides support for parts of this account in generalized anxiety, and momentary studies suggest that emotional contrasts in everyday social interaction may help reinforce worry. ^[UW-04] ^[UW-05] ^[UW-07] The evidence does not establish a universal purpose for all worry. A model can illuminate one route without becoming the explanation for every route.

The idea is psychologically precise because it explains a strange preference: a predictably bad internal world can feel safer than a potentially good but changeable one. Calm becomes dangerous not because calm hurts, but because it creates altitude from which one might fall.

This arrangement is expensive. The person pays a negative emotion now to reduce the possibility of a sharper negative emotion later. The insurance premium is collected every day; the insured event may never occur. Yet the policy can feel successful precisely when nothing happens. “I was prepared” is available as an explanation, while the emotional life that bracing displaced is difficult to count.

Positive states may become complicated too. Joy invites the thought of loss. Hope creates something to disappoint. Rest lowers the guard. The issue is not simply an excess of negative prediction. It is a

difficulty permitting emotional movement without trying to flatten the contrast in advance.

Verbal weather over concrete ground

Worry often arrives as language: rapid strings of “what if,” abstract chains, arguments with imaginary futures. Language is one of humanity's great simulation technologies. It lets us solve a winter problem in summer and prepare for a conversation before damaging a relationship.

Its very flexibility permits a peculiar avoidance. Abstract verbal worry can maintain contact with the category of danger while avoiding the concrete detail that would either enable action or reveal uncertainty's limit.

“What if everything collapses?” can run for hours. “The rent is short by this amount, it is due on this date, and these are the three people or institutions I can contact tomorrow” is less grand and more actionable. Conversely, some problems become painfully concrete only to reveal that no further action is presently available. Worry may return to abstraction because actionlessness is hard to bear.

Not all verbal thought is avoidant, and vivid imagery is not automatically therapeutic. Some people are overwhelmed by concrete catastrophic images. The useful distinction is functional: does the mode of thinking clarify a solvable problem, support tolerable emotional processing, or keep the threat active without creating either action or acceptance?

The answer can change within the same ten minutes. Planning becomes rehearsal; rehearsal becomes rumination; rumination discovers one actionable detail. Mental categories are not separate rooms. The skill lies partly in noticing when the work has changed.

Intolerance of uncertainty, carefully stated

Intolerance of uncertainty is among the best-studied constructs in the worry literature. The phrase does not mean disliking uncertainty; almost everyone dislikes some forms of it. It refers to a tendency for

uncertain situations or thoughts to evoke disproportionately negative interpretation, emotion, and action.

A preregistered 2026 meta-analytic structural study assembled 738 effects from 115 studies or research works, representing 138 independent samples and 28,693 participants. The association between intolerance of uncertainty and worry was substantial, around $r=.59$. It weakened when depression and overlapping questionnaire content were considered, and the evidence remained largely correlational. ^[UW-01] These details matter. “Strong association” is not “single cause,” and a relation between self-report measures is not proof of a neural mechanism.

Earlier meta-analytic work also found that intolerance of uncertainty extends beyond generalized anxiety to depression and obsessive-compulsive symptoms. ^[UW-02] Transdiagnostic does not mean universal, identical, or specific. It means that a process may participate in several different architectures.

The construct is most useful when converted from label to question. What does uncertainty cost this person? Does it imply danger, irresponsibility, shame, helplessness, or endless work? Does the person seek more information, postpone action, demand reassurance, or mentally rehearse? Does the reaction occur broadly or in one domain? A scale score cannot answer all of that.

Productive people in an uninhabitable future

Chronic worry is not always visible as paralysis. It can produce calendars, lists, immaculate preparation, rapid replies, and impressive professional performance. The person may look like the opposite of anxious.

Internally, however, the future has become administrative. Tomorrow is not somewhere to live; it is a department to manage. Each pleasure arrives with its maintenance schedule. Each relationship contains a risk register. Each unfinished task emits a low electrical hum.

High function can therefore conceal high cost. Output tells us that the system moves; it does not tell us how much freedom remains in the movement. A person may complete everything and be unable to stop. They may succeed because the factory never closes, then interpret exhaustion as the price of remaining safe and competent.

This is not an argument against conscientiousness. It is an argument for measuring more than productivity. Can attention leave the problem? Can rest occur without an internal trial? Can an uncertain future contain invitation as well as obligation? Can planning end when the plan is adequate rather than when possibility itself has been defeated?

The process beneath the topics

Content reassurance addresses a product. Chronic worry may be organized around the factory.

Answer the medical question and the system may experience brief relief before locating another unresolved future. This does not make reassurance cruel or useless. Real questions deserve answers. The problem arises when each answer is expected to confer a general certainty no answer can provide.

A formulation therefore asks what rule keeps production active. Perhaps worry is believed to prevent surprise. Perhaps it protects a moral identity. Perhaps abstraction avoids emotion. Perhaps bodily arousal is being interpreted as proof that the question remains urgent. Perhaps the person's life contains ongoing instability that genuinely generates new problems faster than they can be solved.

Those routes are not equivalent. An intervention appropriate for a self-reinforcing certainty ritual may be insulting in a housing crisis. A person can have a worry disorder and real reasons to worry. Mechanism and context must meet.

The endless future of generalized worry is paradoxically narrow: countless threats, few inhabitable destinations. The next chapter turns to threat from the body and from other minds.

Chapter 22 - Threat From Within, Threat From Other Minds

The same room, two different worlds

Two people enter the same meeting. One thinks, “What if my heart accelerates and I cannot get out?” The other thinks, “What if they see that I do not belong here?”

Same chairs, same fluorescent light, same mildly disappointing biscuits. Different universe.

Anxiety organizes a setting according to the predicted location of threat. In panic and some forms of health anxiety, physiology can become the uncertain object. In social anxiety, another person's evaluation may be the critical variable. That variable is fundamentally hidden. No one can inspect another mind directly, which leaves interpretation doing a great deal of work.

These threat locations can combine. A trembling hand may be frightening because it is interpreted as illness, loss of control, or visible evidence of inadequacy. The sensation is the same; the forecast attached to it is not. Mechanisms acquire meaning inside an architecture.

This is why phrases such as “body anxiety” or “fear of judgment” are beginnings, not explanations. We need to know which outcome the person predicts, what they monitor, what they do, and what their actions prevent them from learning.

A body is a data stream that never logs off

Heart rate varies. Breathing changes with posture, exertion, temperature, speech, illness, medication, caffeine, emotion, and attention. Digestion improvises. Muscles twitch. Pain appears, moves, and sometimes leaves without explaining itself.

The body is not a quiet machine that sends a message only when something is wrong. It is a living process full of normal variation. It is also capable of serious illness. The interpretive problem cannot be

solved by declaring every sensation benign or every sensation significant.

Anxiety can alter attention to bodily data, the meaning assigned to it, confidence in that meaning, and the action it triggers. Importantly, a large meta-analysis did not find that anxiety is generally associated with superior objective accuracy at detecting cardiac signals. ^[IP-01] Feeling preoccupied with the heart is not the same as measuring it more accurately. Interoception includes several separable processes: attention, objective performance on a task, subjective sensibility, confidence, and metacognitive fit between confidence and performance.

The distinction is humane. “You are imagining it” is often false-the sensation is real. “Your body is warning you accurately” may also be unwarranted. The central question is how a genuine sensation becomes evidence for a particular cause and a particular action.

Medical safety must remain outside any rhetorical contest. New or severe chest pain, fainting, major shortness of breath, new neurological signs, confusion, or an abrupt change associated with medication or substance use warrants appropriate professional assessment. A popular-science model cannot triage an individual from prose. Refusing false reassurance is part of taking anxiety seriously.

Panic's short clock

Panic compresses time. A sensation that might be observed over minutes becomes a verdict required now. The question is no longer “What pattern will this follow?” but “How do I stop it before the next ten seconds?”

This short clock changes evidence. The natural rise and fall of arousal cannot be learned if escape occurs at the first steep increase and the decline is credited to escape. A peak that lasts minutes can be forecast as endless because the simulation has no credible scene after the peak.

Panic also turns attention into an amplifier-not necessarily by making the bodily signal physically larger, but by giving each

fluctuation interpretive priority. The person checks whether breathing is normal; checking changes breathing. They monitor whether dizziness is worsening; visual fixation, muscle tension, and rapid breathing can alter the experience. The observer joins the experiment.

None of this makes the episode voluntary. A recursive loop can be involuntary at every step. Nor does it imply that panic is only cognition. Carbon-dioxide challenge studies show heightened panic responses in several diagnostic groups, and rare lesion cases demonstrate that panic cannot be assigned to one obligatory “fear centre.” ^[IP-02] ^[IP-03] The body, learning history, context, and interpretation form a system.

The private theatre of public evaluation

Social anxiety solves a different impossible problem: determining one's exact place in other minds.

Was that smile genuine? Did the pause mean boredom? Did I speak too long? Why did she look at her phone? Ordinary social life is saturated with ambiguous evidence because other people have their own histories, distractions, moods, and private concerns. The data are incomplete even on a good day.

The socially anxious person may therefore perform in two theatres at once. In the external room, they are answering a question. In the internal theatre, an imagined audience scores the voice, hands, posture, timing, face, and evidence of anxiety. Attention moves from conversation to self-monitoring and back, often losing something in transit.

The internal audience can produce the very stiffness it fears. Careful control reduces spontaneity; reduced spontaneity feels like proof of poor social ability. Afterward, memory may be reconstructed from the felt strain rather than from what other people actually did. “It felt awful” becomes “I looked awful,” although those are different claims.

Social threat also carries historical realism. Some people have been mocked, excluded, racialized, rejected, or punished for visible difference. A formulation that treats fear of evaluation as baseless may

repeat the social failure. The question is not whether rejection exists. It is how widely, rigidly, and prospectively one history is governing current encounters-and whether the person has access to discrimination, support, and repair.

Blushing changes meaning when it has an audience

A blush is not usually a medical emergency. In social anxiety it can feel catastrophic because it reveals something. The feared event is not heat in the face but involuntary publication: they will see that I am anxious, and seeing it will alter my standing.

A trembling voice can carry several meanings. In panic, it may signal loss of bodily control. In health anxiety, illness. In social anxiety, exposure. In trauma-related states, the beginning of a defensive response associated with an earlier danger.

This is why transdiagnostic physiology must be handled carefully. The same arousal system can participate in distinct predictions and behaviours. A treatment or explanation that targets “the sensation” without asking what the sensation means may miss the controlling model.

The reverse is also true. Different signals can serve the same function. Blushing, a delayed reply, and a forgotten name may all become evidence for social rejection. What looks diverse at the sensory level can converge at the level of predicted consequence.

People regulate people

Relationships are not only possible sources of judgment. Human nervous systems are built in company. A familiar voice, hand, face, or rhythm can change what becomes tolerable.

Small laboratory studies have found that holding the hand of a familiar supportive person can reduce reported distress and alter group-level neural responses during anticipated threat; the effects depend on relationship and context and should not be generalized into a universal biological law. ^[SR-03] ^[SR-04] Other experiments suggest that

cues associated with support can influence fear learning, again in brief nonclinical settings. [SR-07] [SR-08]

The everyday truth is larger than any one imaging result and more conditional. A calm person can lend structure to a chaotic moment. A trusted companion can make a first journey possible. A therapist's steadiness can allow contact with uncertainty that would otherwise be intolerable.

Support and reassurance are not opposites, and no moral border separates them. The useful distinction concerns function over time. Does the relationship help the person remain present, learn, choose, and gradually carry more of the uncertainty? Or must the other person repeatedly remove uncertainty before life can continue?

Support can widen the field. Accommodation can sometimes narrow it. Yet reducing accommodation without preserving warmth can feel like abandonment, especially when dependence developed through genuine need. Change requires co-regulation and autonomy to be treated as partners rather than enemies.

The quantified social field

Modern communication gives other minds instrument panels.

Read at 10:42. Typing... then nothing. Active now. Viewed by 317 people. Three likes. One person removed a reaction. A message that once disappeared into the postal system now emits a chain of micro-observations.

These measurements promise social knowledge but often remain ambiguous. A read receipt shows that a device registered access, not what the person felt. Response time can reflect rejection, work, sleep, illness, distraction, resentment, or a depleted battery. The metric is precise in one dimension and nearly mute in the dimension the anxious system cares about.

Claims that phones simply cause anxiety are too crude. Patterns of problematic use are associated with distress, but directionality, selection, context, and individual vulnerability matter, and the current

evidence ledger does not support a clean causal statement. Technology did not invent the social nervous system. It increased its sampling rate and lowered the cost of checking.

Cheap checking creates dense learning. Refreshing can produce intermittent rewards: sometimes nothing, sometimes the desired reply. Intermittent schedules are good at preserving behaviour because the next attempt might be the one. The result is a social slot machine whose prize is not popularity but temporary closure.

The other person as information and uncertainty

Other people occupy a paradoxical role. They are among our richest sources of information and among our least controllable sources of uncertainty.

A relationship can correct a threat model through years of reliable response. It can also become the domain in which no amount of reliability purchases certainty about the future. Love is evidence; it is not a guarantee. Trust is not a posterior probability of one. It is a way of acting when another person's freedom prevents certainty.

The same is true of the body. A medical assessment can provide strong evidence. It cannot promise that a body will never change. Health requires appropriate attention and an eventual ability to live without continuous proof of continued health.

Internal threat asks, "Can I trust this sensation?" Social threat asks, "Can I trust my place in that mind?" Relationship threat asks, "Can I trust what continues when I cannot inspect it?" These questions share an uncertainty problem, but the objects, histories, harms, and remedies differ.

The distinction prepares us for a broader diagnostic landscape. Similar loops cross conventional categories, yet diagnoses are not arbitrary labels pasted onto one mechanism. The books in psychiatry's library do speak to one another. They do not become the same book.

Chapter 23 - Narrow Fears, Overlapping Worlds

The library whose books keep moving

Psychiatric classification resembles organizing a library while the books hold conversations across the aisles.

Panic disorder, social anxiety disorder, generalized anxiety disorder, specific phobia, obsessive-compulsive disorder, post-traumatic stress disorder, and illness anxiety occupy different shelves. Then avoidance wanders through several of them. Intrusive thoughts visit more than one. Bodily arousal, reassurance, shame, uncertainty, and sleep disruption refuse to respect the catalogue.

This does not make diagnosis useless. Patterns matter. They guide research, communication, eligibility, prognosis, and treatment choices. A diagnosis can help a person discover that an apparently private chaos has a recognizable form.

But categories are maps, not walls. A maintenance process may cross diagnoses, and two people sharing a diagnosis may be sustained by different processes. Classification asks which pattern is present. Formulation asks how this particular pattern keeps recreating itself in this particular life. We need both questions, and neither should impersonate the other.

A small object can govern a large geography

Specific phobias are sometimes treated as the simple residents of the anxiety neighbourhood: one object, one fear, one avoidance. Their consequences can be anything but small.

Fear of flying can redraw a person's map of the planet. Needle fear can alter vaccination, blood tests, fertility treatment, dental care, and pregnancy. Fear of vomiting can shape food, travel, schools, parenthood, and contact with illness. Fear of dogs can determine streets, parks, friendships, and where a child is permitted to play.

The breadth of the cue and the breadth of impairment are separate variables. A narrow trigger can sit on a major bridge. Closing the bridge changes every destination beyond it.

Phobias also reveal why “just avoid it” is sometimes sensible and sometimes costly. Avoiding an obscure animal one never encounters may have almost no consequence. Avoiding injections can carry medical risk. Mechanism cannot determine importance without values and environment.

Fear generalization is relevant but not uniform. Meta-analyses find broader group-level generalization in anxiety-related conditions, with meaningful heterogeneity across tasks, diagnoses, and measures. ^[LE-03]
^[LE-04] A specific phobia is not necessarily evidence that the person's entire learning system generalizes broadly. The fear may be narrow in stimulus space and vast in consequence space.

Obsessive-compulsive disorder is not “anxiety with extra checking”

Obsessive-compulsive disorder should not be collapsed into anxiety. Obsessions, compulsions, incompleteness, “not just right” experiences, disgust, responsibility, taboo thought content, sensory phenomena, and habit-like repetition can form a distinctive architecture. Anxiety may be prominent, secondary, intermittent, or not the only force at all.

Still, OCD exposes one way uncertainty can become behaviourally expensive. Did I lock the door? Did I contaminate someone? Was that movement deliberate? Does having the thought reveal an intention? The problem may not be believing the feared outcome is likely. It may be finding any remaining doubt morally intolerable.

Checking can close the question briefly while weakening trust in memory or completion. Repetition changes the experience of remembering; the act meant to establish certainty can make the event feel less distinct. Then the next check is recruited to repair uncertainty partly created by checking.

Intolerance of uncertainty is associated with OCD symptoms as well as generalized anxiety and depression, which supports overlap without

proving identity. ^[UW-02] A transdiagnostic process can enter different systems and do different work. In generalized worry it may prolong prospective thinking. In OCD it may interact with inflated responsibility, obsessional meaning, disgust, or incompleteness. The shared word “uncertainty” must not erase those companions.

The distinction has practical stakes. Telling someone with OCD simply to relax or think more realistically can become another invitation to settle the obsession on its own terms. Evidence-based approaches often target the relationship between obsession, compulsion, and uncertainty, not merely the intensity of fear. Treatment specifics belong later; here the point is conceptual restraint.

Trauma begins with a fact the model must not insult

Post-traumatic stress cannot be reduced to anxiety either. Trauma-related predictions carry a fact that some generic models evade: the dangerous world was real.

The person did not merely overestimate what humans can do, what a blast can mean, how quickly safety can disappear, or how a body responds under threat. They learned from an event or environment that violated ordinary assumptions. Hypervigilance, avoidance, startle, and anticipatory fear may be organized around that learning, but PTSD can also include intrusive memory, nightmares, dissociation, shame, guilt, emotional numbing, altered self-experience, anger, grief, and disturbances of trust.

In Israel, after October 7 and during continuing periods of war and alarm, the distinction between remembered danger and current danger cannot be handled as a classroom technicality. A siren may indicate a real, time-sensitive threat. People may have relatives serving, displaced communities, ongoing uncertainty, or direct loss. Asking whether the nervous system is “overreacting” without specifying the present context can become both scientifically empty and morally careless.

Calibration after trauma is not a demand to believe the world is safe. It is the difficult work of discriminating contexts, time scales, and available actions while honouring what was learned. The old rule may still be accurate in some settings and overgeneralized in others. Safety can be partial, temporary, socially distributed, and politically contingent.

Neuroimaging meta-analyses show overlap and divergence across PTSD, social anxiety, and specific phobia, with distributed patterns that cannot support a single-centre or single-diagnosis story. ^[NB-01] Those group-level findings do not tell us whether one survivor is safe, what their symptoms mean, or which treatment they need. They mainly warn against neurological cartoons.

Panic is not a phobia of panic, although it can become one

Panic disorder can acquire phobic qualities: certain places, sensations, or distances from help become cues. But the feared object is often a rapidly changing internal event rather than a stable external thing. The person may fear the attack, its consequences, the inability to escape, or the public meaning of losing control.

Agoraphobic restriction can then spread across buses, queues, bridges, cinemas, traffic, and being alone. These places do not resemble one another physically. They resemble one another in predicted escape difficulty, help availability, or humiliation. Generalization follows meaning rather than shape.

This separates panic from generalized worry even when both involve future threat. Worry may run for hours in verbal simulation. Panic collapses the horizon into the next bodily change. Health anxiety may also monitor the body, but its feared inference can concern undetected illness rather than the immediate catastrophe of arousal.

One mechanism-interoceptive attention-therefore cannot stand in for one diagnosis. Objective heartbeat-detection accuracy, subjective bodily concern, catastrophic interpretation, and fear of arousal are separable. ^[IP-01] ^[IP-04] The diagnostic pattern emerges from how those pieces combine with episodes, avoidance, context, and impairment.

Comorbidity is sometimes a clue

When diagnoses co-occur, the overlap is often treated as administrative inconvenience: extra boxes, untidy samples, difficult statistics. It can also reveal coupling.

Panic plus social anxiety may turn a bodily sensation into feared public evidence. Generalized worry plus health anxiety may give topic-switching prospective thought a bodily destination. PTSD plus panic can link trauma cues, arousal, and fear of arousal. Depression can reduce future reward, making anxious avoidance harder to dislodge. Insomnia can increase next-day vulnerability while anxiety makes the following night harder to enter.

These are possible formulations, not automatic explanations of comorbidity. The same pair of diagnoses can couple differently in different people. One may have panic that predates trauma; another develops panic only in trauma-related contexts. One person's depression follows years of restriction; another's low energy independently reduces the capacity to engage in exposure.

Co-occurrence therefore argues for chronology. What began first? What changes what? Which loop operates only in a particular state? What remains when one symptom cluster improves? A flat list of diagnoses contains less causal information than a careful timeline.

It also argues against blaming treatment failure on the person. If an intervention targets one loop while another repeatedly reactivates it, limited improvement may reflect an incomplete model. Diagnosis remains important; formulation supplies the plumbing between the rooms.

Same label, different engine

Two people can meet criteria for social anxiety disorder. One fears visible panic sensations. Another expects contempt after years of bullying. Another is perfectionistic and compares every spontaneous sentence with an impossible internal script. Another lives in an environment where their accent, race, religion, disability, or gender expression does attract hostile attention.

The label identifies a family resemblance. It does not determine which prediction carries most weight, which safety behaviour preserves it, which supports exist, or which risks are real.

Likewise, two people with generalized anxiety may have different engines: intolerance of uncertainty, beliefs about responsibility, emotional contrast avoidance, sleep loss, ongoing material insecurity, trauma-related vigilance, or a combination. Group-level constructs help organize inquiry. They should not be used to overwrite the individual.

Heterogeneity is sometimes presented as a nuisance preventing clean science. It may instead be the phenomenon. A category can contain several paths to a similar surface pattern, and one path can lead to several categories over a lifetime.

Transdiagnostic does not mean identical

A transdiagnostic idea earns its value by crossing boundaries without demolishing them.

Avoidance, uncertainty, prediction, bodily interpretation, and reinforcement appear across disorders. That does not mean every disorder is “really” avoidance or every treatment should be the same. Shared mechanisms interact with disorder-specific content, developmental history, biology, memory, identity, and environment.

The **Flow Hijacked synthesis** will later make a wager: several anxious and anxiety-related patterns may sometimes converge on reduced *future reachability*-fewer valued states remain practically accessible because threat-biased prediction and protective policy repeatedly narrow the route. This is explicitly a Flow Hijacked hypothesis, not a diagnostic fact.

The wager must face discriminating tests. If a measure of reachable, valued alternatives adds nothing beyond symptom severity, depression, disability, and environmental constraint, the new vocabulary has not earned its keep. If different disorders show different patterns of entry, exit, and generalization, the model must

preserve those differences rather than averaging them into a dramatic picture.

No diagnosis should be annexed to make a synthesis look comprehensive. The purpose of the shared lens is to ask an additional question: beyond what the person feels and fears, which futures have become easy to enter, difficult to leave, or absent from practical choice?

The last chapter of this movement expands the frame once more. Anxiety does not occur in an empty laboratory. It occurs in economies, interfaces, wars, families, climates, institutions, and devices designed to make the next uncertain update nearly free.

Chapter 24 - Ancient Alarm, Twenty-First-Century Bandwidth

The visitor from 1820

Bring someone from 1820 into a modern kitchen. Refrigeration, electric light, clean running water, antibiotics in the cabinet, fruit out of season. Place a phone in their hand and explain that it contains maps of the planet, much of recorded knowledge, instant communication, medical information, music without an orchestra, and photographs of people on the other side of the world.

Then show them what we do with this marvel at 2:13 in the morning.

Refresh.

The twenty-first century has solved uncertainty problems with astonishing success. It has also constructed extraordinarily efficient uncertainty-sampling machines. Information has become abundant; certainty has not.

For questions with stable answers, this is a triumph. When does the train leave? What is the address? What does this word mean? For many anxious questions, search changes the size of the hypothesis space faster than it changes the answer. Is this symptom dangerous? Will this relationship last? Is the market about to fall? Has the political situation crossed a threshold? Each result opens links to further possibilities.

Knowledge is not the villain. The same phone can summon medical help, locate shelter, translate instructions, connect a lonely person, document abuse, or warn of incoming danger. The relevant question is not whether technology is good or bad. It is how a particular interface couples with prediction, reinforcement, attention, and agency.

When checking becomes almost free

Once, checking whether a letter had arrived required walking to the door. Checking a share price required a newspaper. Checking another

person's social response might have required a conversation, with all the risk and richness a conversation contains.

Now one thumb can inspect messages, markets, news, deliveries, symptoms, sleep, heart rate, and social approval before the body has fully emerged from bed.

Lowering response cost changes behaviour. If each check briefly reduces uncertainty, making checking nearly effortless allows the relief loop to run hundreds of times. The individual act is too small to look consequential. The day becomes perforated by them.

Intermittent information makes the loop especially sticky. Usually there is nothing; occasionally there is the message, result, or reassuring number. The uncertain schedule keeps the next refresh attractive. An old learning system meets an interface optimized for return.

This is not evidence that design alone causes an anxiety disorder. People select tools for reasons, symptoms influence use, and most digital behaviours are entangled with work and relationships. Causal claims about social media and anxiety remain difficult, and this book's present evidence ledger does not justify a single-direction verdict. Restraint is more informative than a slogan.

Still, the behavioural architecture is visible. The platform can reduce the cost of a policy already reinforced by relief. It becomes part of the loop.

A quantified body is still a living body

Wearable devices can reveal arrhythmias, support training, document sleep patterns, and help people make informed decisions. They can also expose streams of normal variation that previously passed without comment.

Sleep score: 62. Heart-rate variability: down. Resting heart rate: three beats higher. Oxygen saturation: one low reading at 4:07 a.m. The numbers appear precise, and precision in the numerical sense can be mistaken for certainty in the medical sense.

A consumer metric is produced by hardware, algorithms, missing data, and assumptions. Even when technically accurate, its meaning depends on context and trend. The body is not failing an exam each morning. Yet the score can become an externalized interoceptive signal: instead of asking how one feels, one waits for a device to announce whether sleep was acceptable.

The result can be comic and miserable. A person lies awake worrying about tomorrow's sleep score, thereby changing sleep. Measurement joins the process it measures.

The answer is not universal abandonment of tracking. For some people, data enlarge agency. For others, they become surveillance. A useful test is behavioural: does the metric support a proportionate decision, or does every fluctuation demand interpretation and repeated checking? Does it return the person to life, or require life to wait for the next reading?

Threat consumption without an action channel

Global media can bring a catastrophe occurring thousands of kilometres away into the hand within seconds. This creates moral connection, public accountability, and life-saving information. It also creates an unusual regulatory configuration: high salience, repeated exposure, and little direct ability to act.

The nervous system receives cues optimized for attention-urgent language, images, live counters-while the person's action channel may be no larger than another refresh. Concern remains physiologically active because the event is real or possible; helplessness remains because useful action is unclear.

This is threat consumption without an action channel. It differs from direct danger and from practical preparation. The distinction is not moral indifference. One can care deeply while choosing a form and dose of information that preserves the capacity to work, support others, vote, donate, prepare, or rest.

Repeated exposure can also erase scale. A rare event viewed fifty times feels like fifty events. A global stream supplies a numerator

without an intuitive denominator. Availability is not fabrication-the event occurred-but availability changes the felt estimate of what is about to happen here.

The same feed can be protective during emergencies. Reliable updates, clear instructions, and finite check times reduce ambiguity. The design question is whether information is connected to a decision. An alert that tells a person where to shelter performs a different regulatory function from an endless montage of threat.

A siren is not a metaphor

A responsible theory of anxiety must be able to speak in a world where danger is real.

In Israel, a siren can require immediate movement to a protected space. In a war zone, vigilance can save a life. Climate change, political violence, racism, antisemitism, displacement, economic insecurity, pandemics, and unsafe homes are not cognitive distortions. Any account that begins by promising safety will fail both science and the people living the evidence.

The existence of real threat does not make every form of anticipatory regulation equally useful. A person can follow alerts and maintain a shelter while also spending every quiet hour scanning unofficial rumours. Someone can prepare responsibly for extreme heat while being unable to sleep because every warm day is treated as immediate civilizational collapse. Real danger and self-amplifying loops can coexist.

The line between adaptive concern and capture cannot be drawn by asking whether the threat is real. It requires further questions. Is the response proportional to location, timing, and reliable information? Does it support an available action? Can the system stand down when the immediate condition changes? What is the cumulative physiological and social cost? Which people and institutions are actually responsible for reducing the danger?

This final question matters because individual coping can become a political alibi. Anxiety about housing is not solved solely by breathing

differently when rent is unaffordable. Fear under discrimination is not corrected solely through exposure to hostile environments. Climate distress should not be individualized while institutions continue producing the threat. Regulation and collective action belong in the same ecology.

Platforms, workplaces, and the design of uncertainty

The anxious system includes more than a nervous system.

A workplace that sends late-night messages without clear expectations manufactures temporal uncertainty. A health portal that releases abnormal results before explanation can create an interpretive vacuum. An immigration process with indefinite waits converts life plans into suspended variables. Precarious schedules make planning rationally difficult. Notification design decides how often unresolved questions re-enter attention.

These conditions do not affect everyone equally. Money can purchase buffers: time, private space, second opinions, transport, childcare, legal help. Social connection can distribute uncertainty. Power can determine whether one has to wait or makes others wait.

An individual model that ignores these facts may call a person inflexible when the environment is genuinely unforgiving. An environmental model that ignores learning may miss how a once necessary policy continues after conditions improve. The unit of explanation is the coupling.

Institutions can reduce anxiety not only through therapy but through reliable communication, predictable procedures, humane margins for error, and visible routes for repair. Good design does not eliminate uncertainty. It prevents avoidable ambiguity from being used as a tax on attention.

More data, less wisdom

Modern life converts more of the future into numbers: risk scores, forecasts, rankings, projections, fertility estimates, credit scores, performance dashboards. Measurement can reveal patterns

inaccessible to intuition. It can also invite a category error: if a future has a number, it feels as though the uncertainty has been domesticated.

Every forecast contains a model, a population, assumptions, and an error range. A probability is not a promise delivered in decimal form. Nor is a personal metric necessarily a personal destiny. Group-level risk does not tell an individual story without context.

Anxiety can turn data into an infinite staircase. A result answers one question and raises another about accuracy, timing, calibration, or the hidden variable not measured. More information cannot satisfy a stopping rule that demands immunity from surprise.

Wisdom is partly the capacity to connect information to a decision. What would I do differently if the number changed? How often could that change occur? Which source is reliable? When is further sampling likely to improve the choice? When is it feeding the demand that the future sign a guarantee?

These questions are not anti-scientific. They are what science looks like after the graph leaves the paper and enters a life.

Ancient machinery, new bandwidth

The twenty-first century did not invent prediction, social comparison, bodily monitoring, avoidance, or reassurance. It altered their bandwidth.

An old defensive system can now receive thousands of micro-updates per day. It can inspect distant dangers before breakfast, measure sleep before noticing wakefulness, compare a private life with curated public lives, and seek another hypothesis whenever the current one becomes tolerable.

Yet the same century offers unprecedented routes to knowledge, community, treatment, accessibility, and collective coordination. The device that hosts the loop can host the exit. Technology is not an external toxin applied to an innocent nervous system. It is an environment built by nervous systems and then inhabited by them.

The **Flow Hijacked synthesis** that follows will propose that anxiety sometimes becomes a dynamical regime in which certain threat states and protective policies become unusually easy to enter, while valued alternatives become harder to reach. The digital and institutional environment can deepen or loosen that regime by changing sampling cost, reinforcement timing, social regulation, and real opportunity. This is a hypothesis to operationalize, not a scientific verdict on modernity.

Everything needed to reach that proposal already has an independent literature: fear learning, generalization, extinction, avoidance, uncertainty, interoception, prediction, causal inference, reinforcement, exploration, social regulation, and environmental constraint. Flow Hijacked does not discover those mechanisms. Its burden is to assemble them in a way that makes a new, risky prediction.

The central question is now ready to change. We have asked how the future comes to govern the present. We can finally ask what kind of system forms when the same forecasts, actions, and absences of evidence begin to hold one another in place-and what would have to be measured before calling that system real.

PART VI

THE FLOW HIJACKED MODEL

Five proposed constructs, measurable predictions, and ways to be wrong.

Chapter 25 - The Anticipatory Attractor

When tomorrow begins organizing today

Most accounts of anxiety begin with an object: a plane, an illness, a room full of people, an unanswered message. The Flow Hijacked model begins one step earlier, with a change in organization. Nothing decisive has happened, yet the body is preparing, attention is sorting, memory is searching, and the available actions are already being ranked. Tomorrow has not arrived. A representation of tomorrow has.

That distinction matters. The future does not reach backward through time and push the nervous system around. A nervous system capable of prediction creates a present model of what may happen next, and that model can alter the present. A mortgage payment changes spending before it is due. An examination changes sleep before the paper is opened. A feared conversation can occupy a Tuesday even though it is scheduled for Friday. Anticipation is not a ghostly force. It is present physiology, present attention and present action recruited by a represented future.

Sometimes this mobilization is exactly what intelligence requires. A storm forecast makes us bring the chairs indoors. A worrying laboratory result prompts a medical appointment. The problem is not that the organism predicts. The problem begins when a prediction recruits a configuration that helps preserve the prediction whether or not the world continues to support it.

Flow Hijacked calls the first of its five candidate constructs the **Anticipatory Attractor**: a recurrent organism-environment regime in which represented future threat contributes to a self-maintaining defensive state in the present.

The word *candidate* is doing important work. This is a proposed research construct, not a diagnosis, a discovered brain module or a replacement for established accounts of fear learning, uncertainty, attention or avoidance. The scientific literatures reviewed earlier show

group-level differences in safety responding, extinction and fear generalization, alongside substantial variation between tasks and people [LE-01][LE-03][LE-04][LE-08]. Those findings make a dynamical synthesis plausible. They do not prove this one.

A valley is a picture, not an explanation

The familiar image is a marble rolling into a valley. Push it a little and it returns toward the bottom. In mathematics, an attractor is a set toward which nearby trajectories tend to evolve. In life, however, the marble is also remembering, breathing, checking its phone, speaking to another person and changing the terrain as it moves. Human anxiety is not a ball in neural soil.

The landscape image earns its place only if it leads to observations that ordinary symptom totals miss. Let the state of a person in a context at time t be represented by a vector:

$$\mathbf{x}(t) = [a(t), b(t), p(t), u(t), m(t), r(t), \dots]$$

Here, a might index arousal, b bodily sensation, p threat prediction, u experienced uncertainty, m memory activation and r readiness for a defensive response. The coordinates are deliberately provisional. Different questions would require different measurements. The point is not that anxiety can already be reduced to six numbers. The point is that it is a configuration, not a single quantity.

A schematic dynamical model can then be written as:

$$\begin{aligned} \frac{d\mathbf{x}}{dt} = & \mathbf{f}(\mathbf{x}, \boldsymbol{\theta}, t) \\ & + \mathbf{K} \Phi_T(t) \\ & + \mathbf{B} \mathbf{u}(t) \\ & + \boldsymbol{\Sigma}(\mathbf{x}, t) \boldsymbol{\xi}(t) \end{aligned}$$

The function $\mathbf{f}$ represents the system's ordinary coupled dynamics; $\boldsymbol{\theta}$ represents slower conditions such as learning history, sleep debt or continuing stress; Φ_T represents a current forecast of future threat;

$u(t)$ represents inputs from the environment or treatment; and $\xi(t)$ represents fluctuation. This is a scaffold for asking questions, not a literal equation of an anxious brain.

One possible forecast term is:

$$\Phi_T(t) = \mathbb{E}[L(s_{t+\tau}) \mid \mathbf{x}_t] + \lambda_U U_t$$

Expected loss and the cost assigned to uncertainty both contribute to control pressure now. The equation forces a useful question: after present danger and general distress are taken into account, does the represented future measurably alter current state in ways that later make that forecast harder to revise?

Recruitment, not merely repetition

A frightening thought is not an attractor. It may appear, upset us and pass. An attractor-like regime recruits other processes that make return or persistence more likely.

Imagine that an email arrives from a manager: “Can we speak tomorrow?” Threat expectancy rises. The shoulders tighten. Attention gives promotion to memories of mistakes and demotes evidence that the manager often writes brief messages. The person rehearses explanations, checks old correspondence and asks a colleague what the message “really means.” Each act produces some immediate regulation. It also keeps the email at the centre of the day. By evening, exhaustion and narrowed evidence make the original forecast feel even more credible.

No single link in that sequence is novel. Cognitive models describe interpretation, learning models describe avoidance and reinforcement, interoceptive accounts describe the influence of bodily signals, and uncertainty research describes the cost of not knowing. The proposed construct concerns their coupling. Prediction changes the state. The state changes what is sampled. Sampling changes action. Action

changes what can be learned. The represented future gains help from the present it has organized.

This is also why two people with the same weekly symptom score may have different problems. One has sharp surges that settle quickly. The other enters a moderate defensive state and remains there for hours, even after the trigger changes. Their means can be identical. Their transition probabilities, recovery times and vulnerability to re-entry are not.

From ordinary language to measurement

In plain language, the Anticipatory Attractor means: *once the anxious configuration gets going, several parts of the system begin helping it continue or return*. A measurable definition must be less forgiving.

A study would need repeated observations, not one questionnaire. It might measure forecasted threat, bodily arousal, attention, belief updating and action selection before, during and after graded uncertainty. The candidate construct would be supported if a defensive configuration showed at least some of the following:

1. **Persistence:** the joint state lasts longer than the initiating input.
2. **Return:** after a small perturbation, the trajectory moves back toward the defensive configuration.
3. **State dependence:** the same input has different effects depending on the person's immediately preceding state.
4. **Recruitment:** change in one component predicts later change in others, rather than all components merely tracking a common rating.
5. **Transition structure:** entries and exits cluster around reproducible conditions rather than occurring as unstructured noise.
6. **Added prediction:** these dynamic features forecast impairment, recurrence or treatment response beyond symptom severity.

None is established as a clinical marker. They describe what the words would have to mean in data.

Predictions, rivals and a way to lose

The construct predicts that dense within-person trajectories will sometimes outperform average anxiety scores. It predicts that recovery time after a standardized perturbation will matter, that recent history will change response to the same challenge, and that interventions may improve function by altering return dynamics before they greatly change mean distress.

Several simpler explanations could produce the same appearance. Stable trait anxiety might account for persistent high ratings without any attractor. Depression or sleep deprivation might slow recovery. A genuinely threatening environment might keep reactivating defence. Repeated measurement itself might create apparent loops. A person could also be following a deliberate rule rather than being pulled toward a dynamic regime.

The Anticipatory Attractor should therefore be weakened or abandoned if trajectory measures are unreliable; if persistence disappears after current threat, general distress and environmental constraint are controlled; if no return dynamics can be demonstrated; or if the construct adds no out-of-sample prediction beyond simpler models. A label that survives only because every outcome can be redescribed as support is not a scientific construct. It is a decorative synonym.

The humane advantage of dynamics

Dynamical language can sound impersonal, but it protects a crucial distinction: a person can be caught in a pattern without *being* the pattern. Understanding the pattern may not be enough to leave it, because explicit understanding is one variable inside a coupled system. This is not an exemption from agency. It is a more accurate location of agency.

The question changes from “Why do you keep doing this?” to “What makes this transition so likely, what keeps the state active, and what would give another action a fair chance?” That question is kinder because it is more precise. It also leads directly to the next problem. A

defensive state does not remain private inside the body. It selects actions, and selected actions determine which futures can still be entered.

Chapter 26 - How the Future Becomes Smaller

A full calendar can conceal a narrow life

The future is too large to erase. Anxiety does not need to erase it. It can make most routes into it too costly to select.

Suppose tomorrow contains ten physically possible actions: travel, call, apply, exercise, attend, wait, speak, rest, ask and explore. All ten remain available in a dictionary. Yet if seven carry intolerable predicted danger, the lived future contains three. The others have not ceased to exist. They have lost the competition for action.

This distinction between *possible* and *reachable* is central. A city may appear on a map even when every road a particular traveller can use is closed. A valued state may remain imaginable while no admissible action trajectory is selected under current conditions. Anxiety can leave the map intact and remove the roads.

Three of the five Flow Hijacked constructs describe this narrowing: **Defensive Policy Lock-In**, **Corrective Evidence Starvation**, and **Future Reachability Compression**. They are presented together because each changes the meaning of the others. Repeated defensive selection matters partly because it changes the information available for learning; altered learning matters partly because it changes which valued futures remain actionable.

Defensive Policy Lock-In

In everyday speech, a policy is simply a way of acting: leave, check, ask, suppress, rehearse, scan, postpone, carry protection, remain close to an exit. A defensive policy is not automatically irrational. Leaving a burning building is excellent policy. Asking a clinician about a new severe symptom is excellent policy. The scientific question is whether protection continues to dominate when several viable actions exist and the conditions no longer justify such a narrow repertoire.

Let $q(\pi|x_t)$ represent the probability of selecting action policy π in state x_t , and let D be the set of defensive policies relevant to the task. Defensive probability mass is:

$$D_t = \sum_{\pi \in \mathcal{D}} q(\pi | \mathbf{x}_t)$$

The diversity of the whole policy distribution can be summarized, imperfectly, by entropy:

$$H_{\pi}(t) = - \sum_{\pi} q(\pi | \mathbf{x}_t) \log q(\pi | \mathbf{x}_t)$$

Defensive Policy Lock-In is the proposed condition in which threat produces a persistent concentration of selection probability in a narrow protective repertoire, across contexts where multiple workable policies remain available.

Lock-in is not the same as habit. Worry can be linguistically elaborate. Reassurance seeking can involve research, argument and subtle social negotiation. A person may know five alternatives and still find that only the defensive one becomes selectable after arousal crosses a threshold. The defining feature is competition, not simplicity.

One measurable definition would combine a high D_t , reduced policy diversity, persistence across repeated choice opportunities and a disproportionate reduction in valued non-defensive actions as threat rises. Ecological reports could be paired with behavioural tasks and, where justified and consented to, passive measures. The target is not maximal randomness. Choosing freely does not mean choosing every option equally. The relevant loss is flexible access to actions that fit the person's goals and the actual context.

The construct predicts that some people with similar distress will differ in how rapidly their repertoire collapses, and that this collapse will forecast functional restriction. It predicts that lowering distress without changing selection may leave the system vulnerable, while increasing policy flexibility may precede full symptom remission.

Its rivals include ordinary habit strength, lack of resources, fatigue, depressed motivation, explicit preference and genuine danger. It should fail if defensive concentration vanishes after those factors are measured; if it does not generalize beyond the task used to define it; or if a standard symptom score predicts later functioning just as well.

Laboratory work already suggests that avoidance can preserve threat expectations and that some people with anxiety disorders continue costly avoidance despite competing rewards, but the findings are heterogeneous and cannot be turned into a trait shared by everyone [LE-06][LE-07]. The construct is a proposed extension of that evidence, not a new name for it.

Corrective Evidence Starvation

Actions do not merely express beliefs. They build the next dataset.

The person who avoids the elevator receives no elevator journey. The person who enters only while gripping a particular object receives evidence about an elevator-plus-protection condition. The person who checks the door for the fifth time learns that certainty arrived after five checks, not what happens after one. The person who asks for reassurance before opening an email learns about opening an email after reassurance.

This is easy to moralize and important not to. A protective action may be necessary, compassionate or temporarily enabling. The question is not whether support was “bad.” It is what the experience allowed the system to observe.

Corrective Evidence Starvation is the proposed loss of potentially model-revising information caused by the policies that are selected. If $I(o;M)$ represents how informative an observation o would be about the current threat model M , a counterfactual difference can be written schematically as:

$$CES_t = \mathbb{E}_{\pi \sim q_{\text{alt}}}[\mathcal{I}(o_{t+1}; M_t)] \\ - \mathbb{E}_{\pi \sim q_{\text{sel}}}[\mathcal{I}(o_{t+1}; M_t)]$$

The first term is the information expected under viable alternative policies; the second is the information expected under the selected policy. The quantity is difficult because the first term is counterfactual. A study would have to estimate it through randomization, carefully matched alternatives, generative modelling or repeated within-person comparisons. A self-report item asking “Did you avoid evidence?” would not be enough.

The word *corrective* also needs restraint. Evidence is not corrective merely because it is pleasant. If the elevator stops, the threat model may deserve strengthening. If a feared outcome does not occur, the relevant observation may still fail to update belief because the person attributes safety to luck or protection. Evidence becomes corrective only relative to a specified prediction and a plausible alternative explanation.

The construct predicts that policies producing the same immediate relief can differ in later learning because they expose the system to different observations. It predicts that explicit anxiety reduction during an experience is not required for learning, and that a successful exposure can feel difficult while still generating discriminating evidence. It also predicts that some apparent treatment gains will remain tied to the clinic, the therapist or a ritual if those conditions are encoded as necessary causes of safety.

Simpler alternatives include memory bias, weak consolidation, misunderstanding the task, or a threat belief so vague that no observation could disconfirm it. Corrective Evidence Starvation should be weakened if altering information access does not alter later prediction once exposure amount is held constant; if measured evidence access adds nothing beyond avoidance frequency; or if the threat model updates normally despite the proposed starvation.

Experimental work shows that safety behaviour can preserve threat expectations in some conditions, but it does not show that every safety behaviour is harmful in every phase of care ^[LE-06]. That boundary is essential. The aim is not to strip away support for the sake of purity. It is to learn whether support is functioning as a bridge into a wider

evidence field or as a permanent explanation for why catastrophe did not occur.

Future Reachability Compression

Control theory asks whether a system can reach a target state through an admissible sequence of inputs. Human futures are not engineering targets, and meaning cannot be reduced to a state vector. Still, the distinction is illuminating.

Let $\mathcal{S}_V$ be a set of future states the person values. A future state is provisionally reachable if at least one feasible policy gives it more than a minimal probability of access under current conditions:

$$\mathcal{R}_t(\varepsilon) = \left\{ \begin{array}{l} s_{t+\tau} \in \mathcal{S}_V : \\ \exists \pi \in \Pi_{\text{adm}} \text{ such that} \\ P(s_{t+\tau} \mid \mathbf{x}_t, \pi) > \varepsilon \end{array} \right\}$$

A value-weighted size can then be defined as:

$$V_t = \int_{s \in \mathcal{R}_t(\varepsilon)} w(s) d\mu(s)$$

Here $w(s)$ represents personal value and μ a measure over states. No clinic can currently calculate a person's true future volume. The formalism clarifies the proposed object: not the number of items in a diary, but the range and value of trajectories that remain actionable.

Future Reachability Compression is a threat-linked reduction in that effective valued set. A person may work, parent, shop and attend appointments while living inside one rigid corridor: the same route, the same people, the same protective conditions, no unplanned bodily arousal, no unanswered question. Busyness can coexist with severe compression.

A measurable approximation might combine the number of valued domains entered, diversity of routes, context portability, the probability of approach under graded uncertainty and the cost

required to make approach selectable. A good measure would distinguish “I did not choose it because it does not matter to me” from “It matters deeply, but no route felt enterable.” It would also distinguish lack of money, time or safety from anxiety-linked restriction.


The construct predicts that functional recovery can begin as expansion before anxiety disappears. One route becomes available, then another. It predicts that reachability may explain why modest symptom changes sometimes accompany profound life change, and why impressive symptom change sometimes leaves a person's world almost unchanged.

Alternative explanations are substantial: poverty, disability, discrimination, caregiving, depression, physical illness and real political or interpersonal danger can all narrow futures. A model that attributes structural constraint to anxiety would be not merely wrong but unjust. Future Reachability Compression should fail if it cannot distinguish those conditions; if valued repertoire does not change with threat within the same person; or if it adds no useful information beyond established disability and participation measures.

One loop, three distinct claims

The three constructs should not be collapsed into one poetic sentence. Lock-in concerns the distribution of selected actions. Starvation concerns the information those actions make available. Compression concerns the valued futures that remain reachable. They may correlate without being identical.

The proposed maintenance loop is:



Every local move can make sense. A city can optimize one junction and worsen traffic across the network. A company can protect the quarter and destroy a decade of capability. A nervous system can reduce uncertainty now and make tomorrow harder to enter. The system-level cost is not proof that the person made a foolish choice. Often the choice was the best available move in a constrained moment. The scientific task is to discover how more moves become available.

Chapter 27 - Anxious Hysteresis

The way out may not mirror the way in

A small change can start a large defensive episode. Later, reversing that same change may not be enough to end it.

An ambiguous bodily sensation triggers a panic surge. The sensation settles, yet scanning continues. A rumour about layoffs begins three days of alarm. The rumour is corrected, but sleep is already disrupted, mistakes feel more likely, and every new message is read through the activated state. A difficult social encounter ends at nine; the imagined post-mortem continues until two in the morning.

This asymmetry has a name in physics and engineering: hysteresis. A system's current output depends not only on current input but on the path by which the system arrived there. Flow Hijacked calls its fifth candidate construct **Anxious Hysteresis**: path-dependent persistence in which the conditions required to leave a defensive regime differ from the conditions that were sufficient to enter it.

The word is not a claim that anxiety literally obeys magnetic physics. It is a demand for a particular kind of evidence. "Anxiety lasted after the event" is not enough. The proposed signature is a reproducible difference between the activation and recovery paths under comparable inputs.

Two thresholds

Let $z(t)$ represent a graded threat or uncertainty input and $y(t)$ a multidimensional defensive response summarized for a defined task. Suppose the system enters a high-defence regime when increasing input crosses $z_{\uparrow}$, but returns to a lower-defence regime only when input falls below $z_{\downarrow}$.

$$\Delta z_H = z_{\uparrow} - z_{\downarrow}$$

If the paths differ, a loop appears when response is plotted against input. Its area can be summarized as:

$$A_H = \left| \oint y \, dz \right|$$

The loop area is not automatically a biomarker. It can be distorted by fatigue, learning, order effects, measurement lag and simple time-on-task. Yet it makes the claim precise. A larger area means that identical nominal input values are associated with different responses during escalation and recovery.

The bidirectional protocol

A credible test cannot present rising threat alone. It needs a **bidirectional protocol**.

First establish a baseline. Then increase uncertainty or threat in small, ethically safe, individually calibrated steps while repeatedly measuring state, belief, physiology and action. After a defined upper point, decrease the input through the same steps. Counterbalance direction across sessions where possible. Include a neutral control condition, enough time to estimate lag, and preregistered rules for distinguishing hysteresis from habituation, fatigue or sequence effects.

The comparison is not simply “before versus after.” It asks whether the response at, say, level four differs depending on whether the person approached level four from level three or returned to it from level five. Repeating the protocol can test whether the loop is stable, context-dependent or altered by treatment.

Dense data matter. A questionnaire at the beginning and end of a week cannot reveal the shape of a transition. Candidate measures include continuous or frequently sampled arousal, expectancy, attentional capture, escape choice, response latency and recovery after perturbation. No single physiological channel should be treated as anxiety itself.

Early-warning signals, carefully stated

Near some dynamical transitions, systems recover more slowly from small disturbances. Variance may rise. Consecutive observations may become more similar, a pattern called increasing autocorrelation. A trajectory may “flicker” between regimes before settling. These are candidate early-warning signals, not settled clinical facts about anxiety.

For a sampled variable y_t , lag-one autocorrelation can be written as:

$$\rho_1 = \text{corr}(y_t, y_{t-1})$$

Recovery from a small perturbation can be approximated by:

$$y(t) - y^* \approx (y(0) - y^*)e^{-t/\tau_r}$$

where τ_r is a recovery time constant. A rising τ_r , variance or ρ_1 before transition would be interesting. It would not prove that a bifurcation occurred. Trends, circadian rhythms, missing observations and changing environments can produce similar statistics. Any early-warning program must beat those alternatives prospectively, outside the data used to discover it.

The practical hope is modest and important: transition risk might sometimes become visible before a full episode, allowing a person and clinician to alter sleep, context, demands, support or action while the repertoire is still relatively wide. The ethical danger is equally clear. A noisy warning score can itself become an object of checking. A prediction tool that increases vigilance may strengthen the pattern it was meant to detect.

Panic: recovery, reactivation and history

Consider a panic episode on a train. Heat rises, the heart accelerates, the person feels trapped, and a catastrophic interpretation recruits escape. The person leaves at the next stop. Within twenty minutes, physiology settles. Yet the system is not

necessarily back where it began. The platform is scanned. The return journey is postponed. The next heartbeat is interpreted through the memory of the episode.

The following morning, ordinary arousal while climbing stairs may reactivate the configuration. The current bodily input is smaller than the original surge, but its meaning has changed. This is a plausible hysteresis example: the path through yesterday's episode affects today's threshold. It is not proof. Panic can also recur because of conditioning, sleep loss, continuing stress, expectancy, medication or substance effects, or medical conditions. Experimental work shows that panic can be provoked through internal challenges in more than one diagnostic group and even in rare cases without intact bilateral amygdalae, which warns against a single “panic centre” story ^{[IP-02][IP-03]}
[NB-06].

A bidirectional study of panic could induce tolerable bodily sensations under medical and ethical safeguards, increase and then decrease intensity, and test whether catastrophic expectancy, escape choice and recovery trace different paths. A second session could examine reactivation by a weaker cue. The key comparison would be history-sensitive response, not merely peak fear.

What would support or defeat Anxious Hysteresis?

The construct predicts direction-dependent responses to the same graded input, slower recovery near transition, and individual differences in loop area that relate to recurrence or functional restriction. It predicts that treatment can reduce loop area, lower recovery time or widen the range over which flexible policies remain available, even if peak distress is not abolished.

Its competitors include ordinary associative learning, sensitization, depletion, expectancy, demand characteristics and measurement delay. Hysteresis should be rejected if rising and falling paths coincide after lag and order are controlled; if loop estimates are unreliable; if the apparent effect is explained by time alone; or if it offers no prediction beyond recent symptom history.

Five constructs are not six hypotheses

The architecture must now remain clean. The five candidate constructs are:

1. Anticipatory Attractor.
2. Defensive Policy Lock-In.
3. Corrective Evidence Starvation.
4. Future Reachability Compression.
5. Anxious Hysteresis.

They name proposed measurable features of a dynamical account. They are not diagnoses, mechanisms already found in the brain, or six research hypotheses. The following **six testable research hypotheses** are a separate layer. They combine the constructs with adjacent evidence and specify possible empirical programs.

Research hypothesis 1: defensive selection precedes additional narrowing

Within individuals, increases in defensive policy concentration will predict later reduction in valued action diversity beyond current anxiety severity, depression, fatigue, resources and environmental constraint. Evidence against it would include reliable temporal ordering in the opposite direction or no added prediction after covariates. Existing avoidance findings are adjacent support, not a direct test ^{[LE-06][LE-07]}.

Research hypothesis 2: valued futures lose competitive weight

Under anxiety, some valued future states will remain vivid yet exert less influence on choice. This should be measured by separating vividness, specificity, expected value and actual selection. Depression, anhedonia, memory difficulty and verbal ability are major alternatives. The hypothesis fails if anxiety has no relation to value-guided choice after those alternatives are controlled, or if manipulating future representation does not change behaviour.

Research hypothesis 3: support expands exploration only under identifiable conditions

Support from a trusted person will sometimes reduce defensive concentration and increase approach or exploration, not merely reduce reported distress. Familiarity, distraction, demand effects and interpersonal accommodation must be tested. The hypothesis fails if support calms without widening action, or if the effect is fully reproduced by a familiar neutral cue. Laboratory studies show context-dependent acute buffering by support figures, but they do not establish a universal or lasting mechanism ^{[SR-03][SR-04][SR-05][SR-06][SR-07][SR-08]}.

Research hypothesis 4: maintenance has within-person transition structure

Dense time series will reveal reproducible state transitions and direction-dependent recovery in at least a subset of people, beyond stable traits and ordinary autocorrelation. The hypothesis fails if transitions cannot be forecast outside the training sample or if simple habit and recent symptom level explain the data equally well.

Research hypothesis 5: recovery can appear first as valued repertoire expansion

An increase in diverse, personally valued action under uncertainty will sometimes precede large reductions in symptom means and will predict later functioning. Temporary mood elevation, increased money or time, social pressure and rater demand are alternatives. The hypothesis fails if repertoire change has no durability or prospective value once those factors are controlled. This is an outcome proposal, not a new named construct.

Research hypothesis 6: the relevant configuration is distributed

The best-performing model will integrate behaviour, bodily state, context and distributed neural measures rather than depend on a single region or chemical. This follows the heterogeneity of fear and anxiety neurobiology but remains unproven as a Flow Hijacked claim ^{[NB-01][NB-02][NB-03][NB-04][NB-05][NB-06]}. The hypothesis fails if a simpler regional

or non-neural model consistently predicts transitions and outcomes better across tasks, people and independent datasets.

A synthesis must earn its nouns

There is no scientific prize for naming six ideas when five will do. Nor is there merit in treating a borrowed mathematical vocabulary as discovery. The five constructs are useful only if they organize measurements that expose something missed by existing approaches. The six hypotheses are useful only if they can be preregistered, compared with serious alternatives and allowed to fail.

This boundary is part of the model, not a disclaimer attached after the exciting claims. The empirical literature supports many ingredients: learning, avoidance, uncertainty, bodily appraisal, social regulation and treatment effects. The particular architecture proposed here is a Flow Hijacked synthesis. It now owes the reader more than coherence. It owes a research program.

The treatment question can then be stated without magic: if anxiety has narrowed state, action, evidence and future reachability, what kinds of experience allow the system to learn and move differently?

PART VII

MAKING THE FUTURE REACHABLE AGAIN

*Treatment as the recovery of options, not the
manufacture of certainty.*

Chapter 28 - Treatment as New Learning

The nervous system does not revise itself by applause

Advice can be accurate and still arrive in the wrong currency.

“The plane is safe.” “Your heart has been checked.” “Most people will not judge you.” “You can cope.” Such statements may be true, compassionate and necessary. Yet a system shaped by repeated prediction and action does not always revise because another person has supplied a better sentence. It needs observations made under conditions that bear on the feared model.

That is one reason exposure-based treatment is more subtle than “face your fear.” Exposure is not a contest of endurance, a performance of bravery or a method for proving that nothing bad ever happens. It is the deliberate construction of experiences from which new distinctions can be learned.

The older popular story said that fear was a liquid in a tank: remain in the situation until the level falls, and the tank empties. Reduction during a session can be useful, but it is not the whole of learning. Extinction research indicates that new learning can coexist with older threat learning, and that expression of safety may remain sensitive to context. Foundational return-of-fear demonstrations include animal work, while human imaging and clinical findings add related but non-identical evidence; none justify a promise that an old memory has been erased ^{[LE-09][LE-10][NB-04]}.

The more useful question is not only “Did fear go down?” It is “What prediction was tested, what actually happened, what did the person do, and what explanation will the nervous system carry forward?”

Prediction before procedure

An exposure begins intellectually before it begins behaviourally. “Elevators make me anxious” is too broad to test. The predicted catastrophe may be fainting, suffocation, humiliation, loss of control,

being unable to escape, or the belief that anxiety will rise forever. Each prediction requires a different observation.

A simple record can separate four things:

1. **Prediction:** What is expected to happen?
2. **Policy:** What would usually be done to prevent it?
3. **Observation:** What happened, including what was uncomfortable or imperfect?
4. **Update:** What does the observation permit us to say, and what does it not establish?

Suppose a person predicts, “If my heart races and I do not sit down, I will faint.” Walking briskly in a clinically appropriate context may test that prediction more directly than entering a crowded shop while continually checking pulse. If the heart races, dizziness comes and fainting does not, the result is not “bodily sensations are always safe.” It is a narrower and more honest observation: under these conditions, the predicted sequence did not occur.

This precision protects exposure from turning into ideology. Sometimes a feared outcome does occur. A person blushes. Someone notices a tremor. A presentation contains a mistake. The relevant learning may be not “I prevented embarrassment” but “embarrassment was survivable, repair was possible, and the evening continued.” New learning includes coping, variability and recovery, not only non-occurrence.

Violation, variability and discrimination

The informative event is often a discrepancy between prediction and observation. Its rough magnitude can be represented as a prediction error:

$$\delta_t = o_t - \hat{o}_t$$

But a large discrepancy is not automatically useful. If the person explains success entirely by a ritual, luck or the therapist's presence, the observation may leave the central model untouched. Nor should

treatment manufacture maximal surprise without regard to consent or capacity. Learning depends on attention, meaning, attribution, timing and what the person can carry into the next context.

Variability can help prevent safety from becoming tied to one narrow arrangement. Practice can differ in place, time, bodily state and sequence, while remaining collaborative and ethically proportionate. A person who can enter one elevator only at noon with one trusted companion has learned something valuable, but perhaps something conditional. The next question is not to remove every support immediately. It is to discover which elements are bridges and which have become entrance requirements.

Discrimination matters just as much as generalization. Treatment should not teach “approach everything.” It should sharpen the difference between danger, discomfort and uncertainty. Leaving genuine danger is not avoidance pathology. Seeking assessment for a new severe symptom is not capitulation. The aim is a threat system that updates, not one trained to ignore evidence.

Exposure and the three narrowing constructs

In Flow Hijacked terms, exposure can be understood as an intervention on action, information and reachability.

It may loosen **Defensive Policy Lock-In** by making a previously improbable action selectable. It may reduce **Corrective Evidence Starvation** by generating observations unavailable under avoidance. It may reverse **Future Reachability Compression** when the action becomes portable enough to reopen valued trajectories.

That framing produces several useful cautions. Mere presence in a feared situation is not sufficient if covert checking, dissociation or constant reassurance prevents engagement with the relevant prediction. Conversely, a support or safety behaviour is not automatically a contaminant. A graded support can enable entry into a situation that would otherwise remain entirely inaccessible. Experimental evidence shows that safety behaviour can preserve

threat beliefs under some conditions, not that every form of support is harmful ^[LE-06]. Function must be assessed, not assumed.

It also means that success should not be measured only by the lowest fear rating. A person may complete an exposure while highly anxious and nevertheless discover that an action remained possible. Another may feel calm only because the feared condition was subtly excluded. Distress and learning are related, but neither is a transparent proxy for the other.

The therapist as part of the context

Nervous systems do not learn alone. A therapist's pacing, attention, voice and willingness to remain present can alter what becomes possible. Laboratory studies suggest that reminders or contact with trusted people can reduce acute distress or threat-related responding in some contexts, with effects that vary by relationship, age and task ^{[SR-03][SR-04][SR-05][SR-06][SR-07][SR-08]}. These studies do not prove a universal mechanism of co-regulation, but they make the social context impossible to treat as background scenery.

The therapist may initially function as part of a context in which exploration becomes possible. That is not a flaw. A climbing rope is useful precisely because it permits a movement that could not safely be practised otherwise. The problem begins only if the lesson becomes “I can do this *because this particular person certifies safety*.” Then the treatment room may become a greenhouse: life is possible there, but the evidence says little about the weather outside.

The solution is not emotional withdrawal. It is transfer. The therapist and the person can notice what the relationship makes possible, then vary context, widen self-directed choice, invite other forms of support, and test what travels. The therapist lends stability without claiming omniscience. They do not say, “I guarantee nothing bad will happen.” They say, in effect, “We can look together, distinguish danger from alarm, and learn what you can do when certainty is unavailable.”

This protects against **safety-certificate dependence**. The phrase describes a clinical risk, not a sixth Flow Hijacked construct. Reassurance from a therapist can become another check if every action requires the therapist to approve the odds. A relationship that expands agency is different from a relationship that becomes the only admissible cause of safety.

The dose is a sequence of learning opportunities

Therapeutic dose is not simply minutes in a room. It includes how often the relevant prediction is activated, whether the person remains engaged long enough to observe an outcome, whether defensive policies change, how varied the contexts are, and whether learning is revisited between sessions. Ten nominal exposures can contain one learning opportunity repeated ten times.

Nor should treatment become a conveyor belt of ever harder tasks. Difficulty has no intrinsic virtue. The next experiment should be chosen for information and value. Sometimes the high-leverage step is smaller: sending one email without rereading it six times; letting the heart rise during a short walk; remaining in a conversation after a word comes out awkwardly; postponing reassurance long enough to observe what the urge does.

CBT has substantial trial support across anxiety-related conditions, but response and remission vary by diagnosis, comparator, definition and follow-up ^{[TX-01][TX-02][TX-03][TX-04][TX-05][TX-06]}. Longer-term benefit is supported more clearly through roughly the first year than beyond it in the available trial literature, and average benefit does not mean guaranteed or permanent change ^[TX-02]. The evidence warrants offering effective treatments. It does not warrant blaming the person when a standard protocol does not fit.

What new learning looks like outside the clinic

The most convincing sign may be ordinary. The person takes the lift without conducting an internal safety trial. They wait for a reply and continue cooking. They notice a skipped heartbeat, remember the

agreed medical boundary, and return attention to the conversation. They speak while uncertain how they will sound.

Fear may still appear. Old learning may return after stress, in a new place or after time has passed. Return does not erase what was learned; it reveals the conditions under which retrieval failed. The response is not “We are back at the beginning.” It is “What did this context ask the system to remember, and what needs to become more portable?”

Treatment as new learning is less cinematic than conquest. It does not end with fear signing a surrender document. It ends with a growing body of observations: I was uncertain and acted; arousal rose and changed; a mistake occurred and repair remained possible; support helped and some capacity travelled with me. The future gains roads because the person has begun driving on them.

Chapter 29 - Acceptance, Metacognition and the Freedom Not to Obey

When control becomes the trap

An anxious mind can become a full-time workshop for removing anxiety. Thoughts are disputed, sensations monitored, uncertainty processed, sleep evaluated and calm measured to see whether calm is working. The project appears therapeutic. It can also keep threat at the centre of administration.

Acceptance-based and metacognitive approaches ask a different family of questions. What if a thought can be present without becoming an order? What if uncertainty is a condition of action rather than a defect to eliminate? What if attention can move without first winning an argument with the mind?

These approaches are not identical. Acceptance and Commitment Therapy emphasizes psychological flexibility, values and willingness. Metacognitive therapy focuses on beliefs about thinking and patterns such as prolonged worry, rumination, monitoring and attempts at mental control. Mindfulness-based interventions include varied practices and programs that train particular qualities of attention and relationship to experience. Combining them into “be present” would erase the differences that matter.

Acceptance is not agreement

Acceptance is often misunderstood as liking distress, approving injustice or abandoning change. In this context it means allowing an internal event to exist long enough that action need not be organized entirely around removing it.

A thought says, “You will fail.” Agreement says the thought is true. Suppression attempts to force it away. Acceptance says: this event is occurring in the mind, and its presence does not settle the next action.

That move can alter policy competition. If anxiety no longer has to be reduced before speaking, then speaking can enter the action set while anxiety remains. The goal is not a serene performance. It is a wider repertoire.

Values give the repertoire direction. Mathematics can count options; it cannot tell a person whether to invest them in family, art, justice, friendship, work, rest or adventure. “Do more things” is not recovery. A person needs enough freedom for chosen directions to matter.

ACT is supported as an effective psychotherapy, but recent synthesis does not show a general superiority over other active, credible psychotherapies ^[TX-14]. That is not a disappointment. A treatment does not need to win a branding contest to help. The responsible conclusion is that acceptance- and values-based work can be useful, with fit, delivery and individual response still important.

Thoughts about thoughts

Metacognition concerns how we relate to thinking and what we believe thinking does. Worry may be treated as preparation, moral responsibility, protection from surprise or proof of caring. At the same time, the person may believe worry is uncontrollable and dangerous. The mind is employed as a guard and feared as an intruder.

Metacognitive work can examine those beliefs experimentally. Did two hours of worry improve the decision, or mainly create the feeling of labour? Does postponing worry make catastrophe more likely? Does monitoring for an unwanted thought reduce its frequency, or keep the search active? The point is not to prove that thinking is useless. Deliberation is indispensable. The question is when thinking generates decisions and when it generates more thinking.

This distinction is especially important because intolerance of uncertainty is strongly associated with worry at the group level, while measurement overlap, depression and heterogeneity complicate causal conclusions ^{[UW-01][UW-02]}. The mind's demand for certainty can feel like evidence that certainty is attainable if only enough cognition is applied. Often it is a policy for regulating the feeling of not knowing.

Metacognitive therapy has promising evidence across anxiety and depression, but the evidence base includes relatively small samples and possible allegiance effects; broad claims of superiority are not established [TX-12]. Again, useful is enough. Scientific maturity begins where marketing ends.

Mindfulness without incense or mythology

Mindfulness can refer to a trait questionnaire, a meditation practice, a therapeutic component or a multiweek program. Those are not interchangeable. At its most practical, it trains the ability to notice an event with enough resolution that immediate reaction is no longer the only option.

The heart accelerates. The first observation is “fast pulse.” The next event may be “danger.” Then comes scanning, checking or escape. Mindful attention does not guarantee a calm heart. It may create a small interval in which sensation, interpretation and policy can be distinguished.

For some people, breath-focused practice is grounding. For others, especially when bodily attention is already frightening or when trauma-related experiences are active, it may intensify distress. Eyes-open practice, external anchors, movement, shorter intervals or another method may fit better. No one should be told that difficulty meditating reveals resistance or moral failure.

Meta-analytic evidence suggests that mindfulness-based interventions outperform no treatment on average across heterogeneous psychiatric samples, but do not show a clear general advantage over established active treatments [TX-13]. This finding supports availability, not mystification. Mindfulness is neither nothing nor a universal “amygdala off switch.”

Defusion and the grammar of command

Anxious language often arrives in the imperative: check, leave, solve, prepare, do not feel this, do not let them see. Defusion practices alter the grammar. “I will faint” becomes “I am noticing the prediction

that I will faint.” The extra words are not magic. They may make the thought observable as an event rather than transparent as reality.

The test is behavioural. Does this change permit a valued action, better discrimination or more flexible attention? If it becomes a ritual repeated until the thought feels distant enough, defusion can be recruited into the control project. Any method can be converted into a safety behaviour when its hidden rule is “I may act only after this removes anxiety.”

Willingness is not forced endurance

Willingness means consent to experience some uncertainty or discomfort in the service of a chosen action. It does not mean remaining in abuse, neglecting medical symptoms or accepting unsafe conditions. The ethical order is: assess danger, clarify value, choose proportionate action, and make room for the internal events that action evokes.

A person living under real threat may need protection, resources and advocacy more than an exercise in acceptance. A person with new chest pain may need medical evaluation, not a lesson in observing sensation. Psychological flexibility includes changing the environment when the environment is the problem.

One family of changes, several routes

Exposure, acceptance, metacognition and mindfulness may sometimes converge on a functional shift without being reducible to one mechanism. The person acts without waiting for certainty. Worry loses its monopoly on preparation. Attention becomes movable. Bodily arousal becomes information rather than an automatic verdict. A valued direction gains probability.

Flow Hijacked does not claim ownership of these processes. Its contribution is a question: which intervention changes which part of the dynamic configuration, for whom, in what context and with what durability? One person may need direct prediction testing. Another may need freedom from the requirement to settle every thought.

Another may need sleep, medication or protection before psychological experiments become learnable.

The right treatment is not the one with the most elegant theory. It is the one that helps this person regain usable movement while respecting safety, evidence, preference and dignity.

Chapter 30 - Changing the Conditions of Change

A mind is not treated outside a body

It is possible to explain anxiety so psychologically that the organism disappears. The person becomes a bundle of beliefs floating above sleep, hormones, pain, medicines, caffeine, work schedules, illness and other people. This is tidy and false.

Biological conditions do not replace learning and meaning. They alter the terrain on which learning occurs. A severely sleep-deprived person may have less room to notice, discriminate and recover. A medication may reduce enough arousal for practice to become possible, create adverse effects that resemble anxiety, or change rapidly enough to complicate interpretation. Pain, endocrine disturbance, autonomic symptoms, substance effects and respiratory or cardiac conditions can overlap with anxiety without being reducible to it.

The systems question is therefore not “Is it biological or psychological?” It is “Which conditions are currently shaping prediction, arousal, action and recovery, and which of them can responsibly be changed?”

Sleep: the night is part of tomorrow's threshold

Sleep and anxiety are associated in both directions at the population level. Sleep disturbance is common across anxiety-related conditions; longitudinal studies suggest that poor sleep can precede later anxiety and that anxiety can precede insomnia, although definitions, methods and confounding limit strong causal conclusions ^{[SL-01][SL-02][SL-03][SL-04][SL-07][SL-08]}. Subjective sleep difficulties can be larger than differences on some objective measures, which does not make the lived difficulty imaginary ^[SL-03].

Experimental sleep loss has increased next-day anxiety in short laboratory work, and trials aimed at improving sleep have, on average, also improved mental-health outcomes, including anxiety, across

heterogeneous interventions ^{[SL-05][SL-06]}. Neither finding means that every anxiety problem is caused by sleep or cured by repairing it.

In the Flow Hijacked model, sleep is best treated as a possible **loop multiplier**, not a sixth construct. Poor sleep can raise baseline arousal, reduce cognitive flexibility, make ambiguity feel costlier and slow recovery. Anxiety can then increase monitoring and worry at night, further altering sleep. The loop is clinically important even when neither side is the original cause.

The practical inquiry is specific: Is there a stable schedule? Is time in bed becoming a prolonged struggle? Are snoring, breathing pauses, restless legs, nightmares, pain, shift work, caregiving, substance use or medication timing relevant? Is “sleep optimization” itself becoming another anxious control system? Sometimes treating insomnia directly helps. Sometimes a medical sleep disorder needs assessment. Sometimes the aim is to stop turning one poor night into a forecast of a ruined life.

Physical activity without the sermon

Movement belongs in the conversation because a body learns through use, arousal and recovery, not because exercise is a moral test. The relevant questions are concrete. Has life become so restricted that ordinary exertion is now unfamiliar? Is increased heart rate avoided because it resembles panic? Does pain, disability, medication, exhaustion, safety or access limit movement? Would a brief, chosen and medically appropriate activity reopen a valued route?

This chapter's evidence ledger does not supply a dedicated synthesis from which to claim a universal exercise effect for anxiety. The responsible response is not to borrow certainty from another literature. It is to keep physical activity in assessment and shared planning, make no promise of cure, and obtain condition-specific guidance when health factors matter. For some people, movement may provide behavioural activation, interoceptive learning, sleep

support or social connection. For others, the first task is rest, pain care, nutrition, accessibility or protection from overtraining.

An action can be small enough to respect the body and large enough to generate information. A walk may reveal that the heart can accelerate and settle. It may also reveal chest pain that needs evaluation. The model never gets to decide in advance which story is true.

Before calling it anxiety

Anxiety is real even when a medical factor contributes to it. A thyroid disorder does not make fear imaginary. Panic disorder does not make new chest pain irrelevant. Several conditions and substances can produce sensations that overlap with anxiety, and comorbidity is common enough that “either medical or psychological” is a poor frame.

The ledger includes evidence of associations between thyroid autoimmunity and anxiety with high heterogeneity, experimental anxiety or panic after large caffeine challenges in susceptible groups, acute psychiatric effects of THC, and somatic overlap between postural tachycardia and anxiety ^{[DD-01][DD-02][DD-03][DD-04]}. Panic disorder has also been associated with later coronary disease in observational cohorts, which is another reason not to dismiss new symptoms reflexively ^[DD-05]. Glucocorticoids and other medications can contribute to psychiatric symptoms in some circumstances ^[DD-07]. None of these sources permits self-diagnosis from a list.

New or severe chest pain, fainting, marked shortness of breath, a new neurological sign, confusion, or suspected intoxication or withdrawal warrants urgent human medical assessment. A sudden change after starting, stopping or changing a medicine, unusually high caffeine intake, cannabis or another substance also deserves clinical review. The purpose of this boundary is not to frighten the reader. It is to prevent a model of anxiety from becoming a reason to ignore the body.

Medication changes probability, not personhood

Medication is sometimes described as a crutch by people who would never call eyeglasses a moral weakness. It is also sometimes described as a precise chemical correction when the actual evidence is more heterogeneous. Both stories are too simple.

Several medications reduce symptoms on average in short-term trials for generalized anxiety or panic, with important differences in acceptability, adverse effects and evidential certainty ^{[TX-08][TX-09][TX-10]}. Population rankings are not prescriptions for an individual. Diagnosis, comorbidity, previous response, interactions, pregnancy, substance use, adverse effects, access and preference all matter.

Medication can alter the conditions of learning. Reduced arousal may make approach possible; sedation or emotional blunting may make some forms of engagement harder; early side effects may resemble the feared sensations; inconsistent use or abrupt change may destabilize the system. These are clinical questions, not rules that can be settled from this book.

Discontinuation deserves particular care. In randomized withdrawal trials among people who had already responded, stopping antidepressant treatment was associated with more recurrence over follow-up than continuing it, but withdrawal effects, enriched samples and limited follow-up complicate interpretation ^[TX-11]. Benzodiazepine evidence does not justify a blanket conclusion of long-term safety or suitability, and dependence and withdrawal require individualized medical planning ^[TX-20]. Do not start, stop, combine or rapidly change psychiatric medication on the basis of a chapter. Decisions and tapering belong with a qualified prescriber who knows the person and the medicine.

Medication is not evidence that the problem was “purely biological,” just as psychotherapy response is not evidence that it was “all in the mind.” Each changes a coupled system through a different entry point.

Neuromodulation: a frontier, not a shortcut

Neuromodulation is an inviting word. It suggests that if the system is stuck, an external intervention might nudge it into another regime. That possibility deserves research. It also attracts premature certainty.

The anxiety literature on transcranial magnetic stimulation and related neurostimulatory approaches is smaller and more heterogeneous than the better-established evidence bases in conditions such as major depression and obsessive-compulsive disorder. A 2019 synthesis found preliminary positive signals, particularly in generalized anxiety and trauma-related samples, but studies and protocols were small and varied ^[TX-19]. A broader 2026 review of difficult-to-treat anxiety found scattered signals across pharmacological, psychotherapeutic and neurostimulatory strategies without a reliable universal hierarchy ^[TX-21].

That is a reason for careful trials, not clinical dismissal and not promotional confidence. Key questions include target, laterality, frequency, dose, sham integrity, diagnosis, comorbidity, durability, adverse effects and whether symptom change translates into valued functioning. A neural signal is not automatically a clinical mechanism; a short-term mean difference is not restored life.

Flow Hijacked makes no claim that its five constructs currently identify a stimulation target. If future work shows that a particular intervention changes recovery time, transition probability or policy flexibility, that would be evidence to evaluate. Until then, the language of attractors should not be used to decorate an unvalidated protocol.

Conditions are levers, not verdicts

Sleep, activity, medical care, medication, nutrition, substances, pain, relationships and material safety can all influence what the system is able to learn. None explains every person. None makes character the diagnosis.

Sometimes the wisest psychological intervention is to improve a biological condition. Sometimes the wisest biological intervention is

paired with new learning so that reduced symptoms become expanded life. Sometimes the environment continues to deliver genuine danger, and the right intervention is not tolerance but protection.

Changing the conditions of change is not avoiding the “real work.” It is part of the real work. But even excellent conditions do not guarantee a predictable response. When treatment stalls, the next move is not to increase blame. It is to improve the model.

Chapter 31 - When Treatment Does Not Work as Expected

Failure is information, but only after we define it

“Treatment resistant” sounds as if resistance belongs inside the person. The therapy was offered; the person failed to yield. The phrase can conceal more than it reveals.

What outcome did not change? Symptoms, sleep, work, panic frequency, avoidance, quality of life or the ability to act while anxious? Was the treatment delivered at an adequate dose and with fidelity to its active procedures? Could the person attend, afford and trust it? Did the formulation match the problem? Was there a medical factor, substance effect, trauma process, depression, neurodevelopmental difference, ongoing danger or relationship pattern that the original model neglected?

A systematic review found wide inconsistency in definitions of treatment resistance in anxiety research; many studies did not define failure clearly, and relatively few required an adequate psychotherapy trial before applying the label ^[TX-16]. Expert consensus can improve practical definitions, but consensus is not proof that one sequence works for everyone ^[TX-17]. The first scientific act is therefore linguistic: replace a verdict with a question.

Not “Why is this patient resistant?”

But “What exactly was tried, what process was expected to change, what changed, and what remained stable?”

Adequate on paper is not adequate in process

A protocol can be completed without its central learning opportunity ever occurring. Exposure sessions may be scheduled, yet the feared prediction is never activated. The person may enter the situation while using a covert ritual that makes the outcome uninformative. Cognitive work may become endless reassurance. Mindfulness may be practised as a test that must lower anxiety. Medication may be prescribed but

not tolerated, affordable or taken consistently. “Twelve sessions completed” is an administrative fact, not a process measure.

Dose includes frequency, duration and repetition, but also contact with the proposed mechanism. Fidelity does not mean mechanical obedience to a manual. It means preserving the function of the treatment while adapting language, pacing and context to the person. An exposure protocol without prediction testing has lost fidelity even if every worksheet is filled in. A values-based treatment that never reaches behaviour has done the same.

This creates a delicate balance. Too little adaptation can make treatment inaccessible. Too much adaptation can remove the ingredient that gave it a chance to work. The solution is not intuition alone. It is explicit process tracking: What was predicted? What action was attempted? Which supports were present? What was observed? What travelled outside the session?

When treatment becomes reassurance

Therapy can accidentally join the anxious policy system.

The person brings a feared possibility. The therapist supplies a probability estimate. Relief arrives. The following week brings another possibility. Soon the therapeutic relationship is doing excellent work as a reassurance service and very little work as a context for learning under uncertainty.

This does not mean reassurance is forbidden. Human beings need information, validation and calm. In a crisis, co-regulation may be the intervention. The question is temporal and functional: does support help the person approach, discriminate and learn, or does action remain impossible until the therapist issues another certificate?

The same problem can occur with measurement. Symptom tracking can reveal patterns. It can also become surveillance, with every fluctuation interpreted as a forecast. Early-warning signals are useful only if they support proportionate action rather than create a new checking ritual.

Comorbidity is not noise around the “real” anxiety

Clinical lives do not respect chapter boundaries. Depression can reduce the value of futures that anxiety makes difficult to enter. Trauma can make a context dangerous in memory and, sometimes, in current reality. Obsessive-compulsive processes can turn reassurance and checking into highly structured rituals. Substance use may regulate arousal in the short term while altering sleep, withdrawal and threat over time. ADHD can affect planning and task persistence. Pain, endocrine illness, autonomic conditions and medication effects can change bodily evidence.

Treating one strand may still help, but the model should not pretend the other strands are nuisance variables. A person may understand exposure and lack the energy to enact it. They may become less panicked but remain trapped by shame. They may sleep better while a violent home keeps the threat system accurately mobilized. They may reduce worry while poverty continues to close routes.

Comorbidity changes sequencing. Sometimes stabilization or protection must precede intensive learning. Sometimes a direct anxiety intervention improves enough functioning to make other work possible. Sometimes two processes need coordinated treatment. No universal algorithm can settle that order without assessment and preference.

Six questions before escalation

When improvement is absent or fragile, a useful review can proceed through six questions.

1. Was the target specified?

“Less anxiety” may be too broad. The target could be fewer panic-driven exits, less time lost to checking, greater ability to travel, more flexible response to bodily arousal or faster recovery after activation. If the target was vague, apparent failure may be measurement failure.

2. Was the intervention actually received?

Attendance is not receipt. Did the person understand the rationale? Could they engage without being overwhelmed or dissociated? Were between-session practices feasible? Was medication taken as agreed, and were adverse effects assessed rather than interpreted as noncompliance?

3. Did the active process occur?

Did exposure generate relevant evidence? Did cognitive work improve discrimination or only produce counterarguments? Did acceptance widen action or become endurance? Did social support transfer capacity or become required permission?

4. What continued to maintain the pattern?

Defensive Policy Lock-In may persist through a different route than the one treated. Overt avoidance may fall while rumination expands. Corrective Evidence Starvation may continue because every success is attributed to medication, the therapist or luck. Sleep loss or ongoing threat may repeatedly reactivate the state.

5. Was the outcome the right one and the time horizon adequate?

Symptom relief, remission, functioning, tolerability and durability are different outcomes. A treatment may help one and not another. Meta-analytic work supports average CBT benefit, but remission definitions vary and many people do not reach remission ^{[TX-01][TX-03]}. Follow-up evidence becomes thinner over longer periods ^[TX-02]. Early improvement can be real without being permanent; delayed improvement can be missed by a short endpoint.

6. What would distinguish the remaining explanations?

If two formulations suggest different next steps, create the smallest safe observation that separates them. If avoidance is primary, a graded approach test may be informative. If exhaustion is primary, increasing challenge may merely generate failure. If a symptom changed after a medication adjustment, clinical review matters more

than another cognitive exercise. Model revision should change what we do, not merely add a paragraph to the notes.

Escalation is not automatically improvement

When a first treatment does not help, the temptation is to add: another medicine, another technique, another device, another diagnosis. Sometimes augmentation is appropriate. But systematic evidence has not established one broadly convincing augmentation strategy for treatment-resistant anxiety, and the available literature is limited and heterogeneous ^{[TX-18][TX-21]}. More treatment can also mean more adverse effects, cost, complexity and uncertainty about what caused change.

A careful next step may be a better-delivered version of an evidence-based treatment, a different established treatment matched to preference, a medical review, a reformulation, an environmental intervention or coordinated care. It may also be a pause to recover from treatment burden. Shared decision-making is not a courtesy added after experts decide. It is part of selecting an intervention that can actually be enacted.

Flow Hijacked as failure analysis

The model offers a set of questions, not a tribunal.

Did the Anticipatory Attractor change in persistence or recovery? Did Defensive Policy Lock-In loosen when threat rose? Did the person gain access to observations that could revise the threat model? Did valued futures become more reachable? Did the activation and recovery paths become less asymmetric?

If none changed, several explanations remain possible. The constructs may have been measured badly. The intervention may not have touched them. Environmental threat may have overwhelmed the experiment. Or the Flow Hijacked model may be wrong for this person or wrong altogether.

That final option must remain live. A model that turns every nonresponse into proof of deeper “hijacking” is unfalsifiable and cruel. The person is not obligated to vindicate the theory.

The dignity of revision

Good treatment contains the same humility as good science. It makes a prediction, intervenes, observes and revises. It does not confuse confidence with fidelity.

A person who has not improved has already supplied important information. Perhaps the treatment targeted arousal when shame governed avoidance. Perhaps it targeted thoughts while the body remained sleep deprived. Perhaps it asked for acceptance in a dangerous environment. Perhaps it was technically sound and simply did not help.

The response is curiosity with consequences. Change the formulation. Reconsider the diagnosis. Check the body. Ask what mattered to the person. Inspect the actual dose. Seek consultation. Preserve what helped. Stop what harms. Try another credible route when the expected benefit justifies its burden.

Treatment failure is painful. It should never be romanticized as a hidden lesson the person needed. Yet it can become usable evidence when it changes the next decision. That is model revision in its most human form: not protecting the elegance of our explanation, but protecting the person's remaining futures.

Chapter 32 - Making the Future Reachable Again

The corridor

We began with a future that arrived early. Before the conversation, the journey, the result or the bodily change, anticipation reorganized the present. Across these chapters, that simple event opened into learning, memory, uncertainty, physiology, attention, action, relationship, culture and time.

Now return to a corridor with many doors.

Some doors may lead to pleasure. Some to boredom. One may lead to pain. None is transparent, and no honest authority can place a sign above them reading **ALL DOORS SAFE**.

The anxious strategy is to wait for enough information. Waiting prevents the immediate risk of opening the wrong door. No catastrophe occurs in the corridor. This makes waiting look wise. Meanwhile, hours become years, and the corridor quietly becomes a life.

Recovery does not require the doors to become transparent. It requires the possibility of opening one before certainty.

What the model claims, and no more

The Flow Hijacked synthesis proposes that anxiety may become especially restrictive when represented future threat helps organize a persistent defensive state; when action selection concentrates in a protective repertoire; when those actions restrict access to model-revising evidence; when valued futures lose actionable routes; and when the path out requires different conditions from the path in.

Its five candidate constructs are the Anticipatory Attractor, Defensive Policy Lock-In, Corrective Evidence Starvation, Future Reachability Compression and Anxious Hysteresis. They are not established clinical entities. They earn a place only if reliable

measures can be built, serious alternatives can be tested, and prediction improves beyond existing models.

The six research hypotheses set out in Chapter 27 are another layer. They propose temporal, behavioural, social, outcome and distributed-system tests. They should not be counted as six more constructs or presented as findings. Terminology is cheap. Prediction is the price of admission.

Recovery as repertoire, not performance

Symptom relief matters. Severe anxiety causes real suffering and deserves effective care. Yet freedom may begin to return unevenly, before a score crosses a remission threshold.

The person still becomes anxious, but more actions remain possible. Recovery after a surge becomes faster. A plan can bend without collapsing. Support is available without becoming mandatory permission. New learning travels into another room, another relationship, another day. The body can object without acquiring final authority.

This is not an argument against symptom measures. It is an argument for adding life to them. How many valued domains remain reachable under uncertainty? How costly is approach? How much of the repertoire disappears as threat rises? How quickly does it return? These questions should complement validated measures, not replace them with a poetic score.

Nor is maximal variety the goal. A meaningful life contains commitments, routines and chosen constraints. The person does not need every door. They need enough access that values can influence which door is opened.

Values choose the direction

No equation can decide that a conversation with a child matters more than a perfectly calm evening, that art is worth possible criticism, or that rest matters more than another hour of checking. Reachability without value is only mobility.

This is why treatment cannot be reduced to lowering a threat parameter. A person may become less anxious and remain directionless. Another may choose a difficult future because it contains love, contribution or dignity. The relevant expansion is not “more behaviour” in the abstract. It is renewed influence of what matters.

Values are not another source of pressure. “If you cared enough, you would do it” is coercion disguised as therapy. A value offers direction; it does not cancel disability, danger or the need for support. Movement can be tiny and still be aligned.

A threat system worth keeping

Fear has a legitimate job. A person with no capacity for threat would not be liberated. They would be vulnerable. Recovery cannot mean teaching the organism that the world is harmless, because it is not.

A good threat system discriminates. It can rise when evidence warrants, fall when conditions change, remain uncertain when evidence is mixed, and share control with curiosity, attachment, play, reward, rest and meaning. It can say “danger” without converting every possibility into a veto.

We should not romanticize anxiety. Suffering is not a gift people are required to appreciate. Anxiety can consume years, restrict bodies, isolate families, derail work and compress identity. The hopeful fact is not that pain was secretly necessary. It is that the capacities anxiety recruits-prediction, care, imagination, foresight, learning-are capacities a person still needs. The aim is not to destroy them. It is to return them to proportion.

The person is not the pattern

A person can be caught in an anxious regime. They are not the regime. A person can avoid repeatedly. They are not avoidance. A person can need medication, therapy, practical support or protection. None becomes their definition.

Systems are histories acting in current conditions. Histories matter, but they are not solitary authors. Bodies change. Environments

change. Relationships change. Policies change. New observations become possible. Sometimes the landscape changes slowly; sometimes a small shift opens a route that had been nearly impossible to select.

This is not a promise that every condition can be changed enough, nor that effort controls outcome. Real constraints remain real. The claim is narrower: identity should not be inferred from the persistence of a state.

The ordinary experiment

Change is often embarrassingly ordinary. One elevator. One unanswered message. One imperfect decision. One walk in which the heart is permitted to accelerate within an agreed medical boundary. One presentation containing a stumble. One evening in which tomorrow is not solved before sleep.

The event does not announce itself as a turning point. There is no orchestra. A trajectory that previously had almost no probability of selection begins to compete.

Someone else may say, “You can cope,” and sometimes that matters enormously. Eventually the system also needs some of its own observations:

I was uncertain.

I moved.

I became afraid.

I stayed long enough to learn something.

I made a mistake.

Repair remained possible.

I did not know.

I chose.

The world did not become certain. Certainty became less necessary for movement.

One ending

A book about anxiety is tempted to end repeatedly. First the science concludes. Then recovery receives a definition. Then the corridor returns. Then courage is praised. Repetition can imitate closure without earning it.

There is only one ending here.

The future remains unfinished. Some dangers are real. Some losses cannot be prevented. Some treatments will disappoint. Some doors should remain closed, and some will open onto rooms we would not have chosen.

What can change is the requirement that uncertainty be resolved before life resumes. Fear can become information rather than government. Support can become a bridge rather than a border. A body can mobilize and recover. A thought can appear without receiving administrative authority. A valued route can become selectable while the outcome is still unknown.

That is not fearlessness. It is not permanent safety. It is a future with roads again.

Epilogue - Before Certainty

Before certainty, there is usually a smaller event.

A hand leaves the door after one check. A message remains unanswered while dinner is made. Someone enters the room with a prepared sentence and discovers that the sentence comes out imperfectly. A train door closes. The heart accelerates. The person notices the old forecast arriving with extraordinary confidence.

Nothing in that moment proves the future safe.

The evidence assembled in this book supports broad roles for threat learning, generalization, avoidance, uncertainty, worry, bodily appraisal, reinforcement, social context and psychological treatment in anxiety-related conditions. It also refuses a single anxious brain, one learning defect, one bodily error or one treatment that explains everyone ^{[LE-01][LE-03][IP-01][UW-01][NB-01][TX-01]}. Heterogeneity is not a footnote to the science. It is part of what the science says.

The Flow Hijacked synthesis adds five candidate constructs and six distinct research hypotheses. It does not convert them into facts by printing them in capital letters. The constructs must become reliable measurements. The hypotheses must face competing explanations. A result that contradicts the model must be permitted to contradict it.

That discipline does not empty the ideas of human meaning. It protects the person from being made responsible for the theory. If future reachability proves useful, measure it. If hysteresis predicts recurrence, test it prospectively. If a supportive relationship widens exploration, discover when and for whom. If symptom reduction and restored life diverge, learn from the divergence. If simpler explanations perform better, use them.

Meanwhile, life does not wait for the final model.

We move with partial evidence, changing bodies, other people, values and histories that do not fit neatly into equations. We distinguish danger where we can. We ask for help. We revise. We

refuse the cruelty of calling every limitation fear, and the equal cruelty of treating every fear as a command.

Courage is sometimes imagined as the feeling that arrives after uncertainty has been defeated. More often it is an action with uncertainty still in the room. The hands may shake. The forecast may remain loud. The next step does not certify a happy ending. It only makes contact with a future that waiting could never reveal.

The task is not to make the future harmless.

It is to make the future reachable again.

Then, somewhere before certainty, we move.

Evidence Notes and References

The compact markers used throughout the book resolve here to the source ledger. A indicates a strong direct synthesis or trial for the stated claim, B a useful but limited or heterogeneous source, and C indirect, small, foundational, or illustrative evidence. A rating never licenses a claim broader than the study measured.

Learning, extinction, and avoidance

- LE-01 · Grade A. Duits P et al. *Updated meta-analysis of classical fear conditioning in the anxiety disorders*. *Depress Anxiety*. 2015;32:239-253. DOI [10.1002/da.22353](https://doi.org/10.1002/da.22353); PMID [25703487](https://pubmed.ncbi.nlm.nih.gov/25703487/)
- LE-02 · Grade B. Lissek S et al. *Classical fear conditioning in the anxiety disorders: a meta-analysis*. *Behav Res Ther*. 2005;43:1391-1424. DOI [10.1016/j.brat.2004.10.007](https://doi.org/10.1016/j.brat.2004.10.007)
- LE-03 · Grade A. Cooper SE et al. *A meta-analysis of conditioned fear generalization in anxiety-related disorders*. *Neuropsychopharmacology*. 2022;47:1652-1661. DOI [10.1038/s41386-022-01332-2](https://doi.org/10.1038/s41386-022-01332-2); PMID [35501429](https://pubmed.ncbi.nlm.nih.gov/35501429/); PMCID [PMC9283469](https://pubmed.ncbi.nlm.nih.gov/PMC9283469/)
- LE-04 · Grade A. Fraunfelter L et al. *Fear one, fear them all: a systematic review and meta-analysis of fear generalization in pathological anxiety*. *Neurosci Biobehav Rev*. 2022;139:104707. DOI [10.1016/j.neubiorev.2022.104707](https://doi.org/10.1016/j.neubiorev.2022.104707); PMID [35643120](https://pubmed.ncbi.nlm.nih.gov/35643120/)
- LE-05 · Grade C. Phelps EA et al. *Extinction learning in humans: role of the amygdala and vmPFC*. *Neuron*. 2004;43:897-905. DOI [10.1016/j.neuron.2004.08.042](https://doi.org/10.1016/j.neuron.2004.08.042)
- LE-06 · Grade B. Lovibond PF et al. *Safety behaviours preserve threat beliefs: protection from extinction of human fear conditioning by an avoidance response*. *Behav Res Ther*. 2009;47:716-720. DOI [10.1016/j.brat.2009.04.013](https://doi.org/10.1016/j.brat.2009.04.013)
- LE-07 · Grade B. Pittig A et al. *Elevated costly avoidance in anxiety disorders: patients show little downregulation of acquired avoidance in face of competing rewards for approach*. *Depress Anxiety*. 2021. DOI [10.1002/da.23119](https://doi.org/10.1002/da.23119)
- LE-08 · Grade B. Duits P et al. *Latent class growth analyses reveal overrepresentation of dysfunctional fear conditioning trajectories in patients with anxiety-related disorders*. *J Anxiety Disord*. 2021;78:102361. DOI [10.1016/j.janxdis.2021.102361](https://doi.org/10.1016/j.janxdis.2021.102361)
- LE-09 · Grade C. Rescorla RA, Heth CD. *Reinstatement of fear to an extinguished conditioned stimulus*. *J Exp Psychol Anim Behav Process*. 1975;1:88-96. DOI [10.1037/0097-7403.1.1.88](https://doi.org/10.1037/0097-7403.1.1.88)
- LE-10 · Grade C. Bouton ME, King DA. *Contextual control of the extinction of conditioned fear: tests for the associative value of the context*. *J Exp Psychol Anim Behav Process*. 1983;9:248-265. DOI [10.1037/0097-7403.9.3.248](https://doi.org/10.1037/0097-7403.9.3.248)

Interoception and panic

- IP-01 · Grade A. Adams KL et al. *Anxiety and cardiac interoceptive accuracy: a systematic review and meta-analysis*. *Neurosci Biobehav Rev*. 2022;140:104754. DOI [10.1016/j.neubiorev.2022.104754](https://doi.org/10.1016/j.neubiorev.2022.104754); PMID [35798125](https://pubmed.ncbi.nlm.nih.gov/35798125/)

- IP-02 · Grade A. Tural U, Iosifescu DV. *A systematic review and network meta-analysis of carbon dioxide provocation in psychiatric disorders*. J Psychiatr Res. 2021;143. DOI [10.1016/j.jpsychires.2020.11.032](https://doi.org/10.1016/j.jpsychires.2020.11.032); PMID [33250190](https://pubmed.ncbi.nlm.nih.gov/33250190/)
- IP-03 · Grade C. Feinstein JS et al. *Fear and panic in humans with bilateral amygdala damage*. Nat Neurosci. 2013;16:270-272. DOI [10.1038/nn.3323](https://doi.org/10.1038/nn.3323); PMID [23377128](https://pubmed.ncbi.nlm.nih.gov/23377128/); PMCID [PMC3739474](https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3739474/)
- IP-04 · Grade B. Olatunji BO, Wolitzky-Taylor KB. *Anxiety sensitivity and the anxiety disorders: a meta-analytic review and synthesis*. Psychol Bull. 2009;135:974-999. DOI [10.1037/a0017428](https://doi.org/10.1037/a0017428); PMID [19883144](https://pubmed.ncbi.nlm.nih.gov/19883144/)
- IP-05 · Grade B. Pompoli A et al. *Dismantling cognitive-behaviour therapy for panic disorder: a systematic review and component network meta-analysis*. Psychol Med. 2018;48:1945-1953. DOI [10.1017/S0033291717003919](https://doi.org/10.1017/S0033291717003919); PMID [29368665](https://pubmed.ncbi.nlm.nih.gov/29368665/)
- IP-06 · Grade C. Hermans D et al. *Priming associations between bodily sensations and catastrophic misinterpretations: specific for panic disorder?* Behav Res Ther. 2010. DOI [10.1016/j.brat.2010.05.015](https://doi.org/10.1016/j.brat.2010.05.015)

Uncertainty and worry

- UW-01 · Grade A. Akbari M et al. *A hierarchical uncertainty framework of anxiety: preregistered meta-analytic structural model of intolerance of uncertainty, worry, anxiety, and depression*. Clin Psychol Rev. 2026;128:102773. DOI [10.1016/j.cpr.2026.102773](https://doi.org/10.1016/j.cpr.2026.102773); PMID [42419154](https://pubmed.ncbi.nlm.nih.gov/42419154/)
- UW-02 · Grade A. Gentes EL, Ruscio AM. *A meta-analysis of the relation of intolerance of uncertainty to symptoms of generalized anxiety disorder, major depressive disorder, and obsessive-compulsive disorder*. Clin Psychol Rev. 2011;31:923-933. DOI [10.1016/j.cpr.2011.05.001](https://doi.org/10.1016/j.cpr.2011.05.001); PMID [21664339](https://pubmed.ncbi.nlm.nih.gov/21664339/)
- UW-03 · Grade B. Ladouceur R et al. *Experimental manipulation of intolerance of uncertainty: a study of a theoretical model of worry*. Behav Res Ther. 2000;38:933-941. DOI [10.1016/S0005-7967\(99\)00133-3](https://doi.org/10.1016/S0005-7967(99)00133-3)
- UW-04 · Grade B. Llera SJ, Newman MG. *Rethinking the role of worry in generalized anxiety disorder: evidence supporting a model of emotional contrast avoidance*. Behav Ther. 2014;45:283-299. DOI [10.1016/j.beth.2013.12.011](https://doi.org/10.1016/j.beth.2013.12.011); PMID [24680226](https://pubmed.ncbi.nlm.nih.gov/24680226/)
- UW-05 · Grade B. Crouch TA et al. *Prospective investigation of emotional contrast avoidance in generalized anxiety disorder*. Behav Ther. 2017;48:544-556. DOI [10.1016/j.beth.2016.10.001](https://doi.org/10.1016/j.beth.2016.10.001); PMID [28577589](https://pubmed.ncbi.nlm.nih.gov/28577589/)
- UW-06 · Grade B. Llera SJ, Newman MG. *Development and validation of two measures of emotional contrast avoidance*. J Anxiety Disord. 2017. DOI [10.1016/j.janxdis.2017.04.008](https://doi.org/10.1016/j.janxdis.2017.04.008); PMID [28500921](https://pubmed.ncbi.nlm.nih.gov/28500921/)
- UW-07 · Grade B. Newman MG et al. *The naturalistic reinforcement of worry from positive and negative emotional contrasts: results from a momentary assessment study within social interactions*. J Anxiety Disord. 2022;92:102634. DOI [10.1016/j.janxdis.2022.102634](https://doi.org/10.1016/j.janxdis.2022.102634); PMID [36182690](https://pubmed.ncbi.nlm.nih.gov/36182690/); PMCID [PMC10187062](https://www.ncbi.nlm.nih.gov/pmc/articles/PMC10187062/)
- UW-08 · Grade B. LaFreniere LS, Newman MG. *Reducing contrast avoidance in GAD by savoring positive emotions: outcome and mediation in a randomized controlled trial*. J Anxiety Disord. 2023;93:102659. DOI [10.1016/j.janxdis.2022.102659](https://doi.org/10.1016/j.janxdis.2022.102659)

UW-09 · Grade A. Songco A et al. *A systematic review of the cognitive model of worry in children and adolescents*. Clin Child Fam Psychol Rev. 2020. DOI [10.1007/s10567-020-00311-7](https://doi.org/10.1007/s10567-020-00311-7); PMID [31989444](https://pubmed.ncbi.nlm.nih.gov/31989444/)

Social regulation and attachment

SR-01 · Grade A. Colonnese C et al. *The relation between insecure attachment and child anxiety: a meta-analytic review*. J Clin Child Adolesc Psychol. 2011;40:630-645. DOI [10.1080/15374416.2011.581623](https://doi.org/10.1080/15374416.2011.581623); PMID [21722034](https://pubmed.ncbi.nlm.nih.gov/21722034/)

SR-02 · Grade A. Groh AM et al. *The significance of insecure and disorganized attachment for children's internalizing symptoms: a meta-analytic study*. Child Dev. 2012;83:591-610. DOI [10.1111/j.1467-8624.2011.01711.x](https://doi.org/10.1111/j.1467-8624.2011.01711.x); PMID [22235928](https://pubmed.ncbi.nlm.nih.gov/22235928/)

SR-03 · Grade C. Coan JA et al. *Lending a hand: social regulation of the neural response to threat*. Psychol Sci. 2006;17:1032-1039. DOI [10.1111/j.1467-9280.2006.01832.x](https://doi.org/10.1111/j.1467-9280.2006.01832.x); PMID [17201784](https://pubmed.ncbi.nlm.nih.gov/17201784/)

SR-04 · Grade B. Coan JA et al. *Relationship status and perceived support in the social regulation of neural responses to threat*. Soc Cogn Affect Neurosci. 2017;12:1574-1583. DOI [10.1093/scan/nsx091](https://doi.org/10.1093/scan/nsx091); PMID [28985422](https://pubmed.ncbi.nlm.nih.gov/28985422/); PMCID [PMC5647795](https://pubmed.ncbi.nlm.nih.gov/pmc/PMC5647795/)

SR-05 · Grade B. Gee DG et al. *Maternal buffering of human amygdala-prefrontal circuitry during childhood but not during adolescence*. Psychol Sci. 2014;25:2067-2078. DOI [10.1177/0956797614550878](https://doi.org/10.1177/0956797614550878); PMID [25280904](https://pubmed.ncbi.nlm.nih.gov/25280904/)

SR-06 · Grade B. Hostinar CE et al. *Parent support is less effective in buffering cortisol stress reactivity for adolescents compared to children*. Dev Sci. 2015;18:281-297. DOI [10.1111/desc.12195](https://doi.org/10.1111/desc.12195); PMID [24942038](https://pubmed.ncbi.nlm.nih.gov/24942038/); PMCID [PMC4270957](https://pubmed.ncbi.nlm.nih.gov/pmc/PMC4270957/)

SR-07 · Grade B. Hornstein EA et al. *A safe haven: investigating social-support figures as prepared safety stimuli*. Psychol Sci. 2016;27:1051-1060. DOI [10.1177/0956797616646580](https://doi.org/10.1177/0956797616646580); PMID [27324266](https://pubmed.ncbi.nlm.nih.gov/27324266/)

SR-08 · Grade B. Hornstein EA, Eisenberger NI. *Unpacking the buffering effect of social support figures: social support attenuates fear acquisition*. PLoS One. 2017;12:e0175891. DOI [10.1371/journal.pone.0175891](https://doi.org/10.1371/journal.pone.0175891); PMID [28463999](https://pubmed.ncbi.nlm.nih.gov/28463999/); PMCID [PMC5413011](https://pubmed.ncbi.nlm.nih.gov/pmc/PMC5413011/)

Sleep and anxiety

SL-01 · Grade B. Alvaro PK et al. *A systematic review assessing bidirectionality between sleep disturbances, anxiety, and depression*. Sleep. 2013;36:1059-1068. DOI [10.5665/sleep.2810](https://doi.org/10.5665/sleep.2810); PMID [23814343](https://pubmed.ncbi.nlm.nih.gov/23814343/); PMCID [PMC3669059](https://pubmed.ncbi.nlm.nih.gov/pmc/PMC3669059/)

SL-02 · Grade A. Cox RC, Olatunji BO. *A systematic review of sleep disturbance in anxiety and related disorders*. J Anxiety Disord. 2016;37:104-129. DOI [10.1016/j.janxdis.2015.12.001](https://doi.org/10.1016/j.janxdis.2015.12.001); PMID [26745517](https://pubmed.ncbi.nlm.nih.gov/26745517/)

SL-03 · Grade A. Cox RC, Olatunji BO. *Sleep in the anxiety-related disorders: a meta-analysis of subjective and objective research*. Sleep Med Rev. 2020;51:101282. DOI [10.1016/j.smrv.2020.101282](https://doi.org/10.1016/j.smrv.2020.101282); PMID [32109832](https://pubmed.ncbi.nlm.nih.gov/32109832/)

SL-04 · Grade A. Hertenstein E et al. *Insomnia as a predictor of mental disorders: a systematic review and meta-analysis*. Sleep Med Rev. 2019;43:96-105. DOI [10.1016/j.smrv.2018.10.006](https://doi.org/10.1016/j.smrv.2018.10.006); PMID [30537570](https://pubmed.ncbi.nlm.nih.gov/30537570/)

SL-05 · Grade B. Ben Simon E et al. *Overanxious and underslept*. Nat Hum Behav. 2020;4:100-110. DOI [10.1038/s41562-019-0754-8](https://doi.org/10.1038/s41562-019-0754-8); PMID [31685950](https://pubmed.ncbi.nlm.nih.gov/31685950/)

- SL-06 · Grade A. Scott AJ et al. *Improving sleep quality leads to better mental health: a meta-analysis of randomised controlled trials*. Sleep Med Rev. 2021;60:101556. DOI [10.1016/j.smrv.2021.101556](https://doi.org/10.1016/j.smrv.2021.101556); PMID [34607184](https://pubmed.ncbi.nlm.nih.gov/34607184/); PMCID [PMC8651630](https://pubmed.ncbi.nlm.nih.gov/PMC8651630/)
- SL-07 · Grade B. Jansson-Fröjmark M, Lindblom K. *A bidirectional relationship between anxiety and depression, and insomnia? A prospective study in the general population*. J Psychosom Res. 2008;64:443-449. DOI [10.1016/j.jpsychores.2007.10.016](https://doi.org/10.1016/j.jpsychores.2007.10.016)
- SL-08 · Grade B. Nguyen et al. *Poor sleep and the development of major depressive disorder and generalized anxiety disorder: an 18-year longitudinal study*. J Anxiety Disord. 2022;90:102601. DOI [10.1016/j.janxdis.2022.102601](https://doi.org/10.1016/j.janxdis.2022.102601)

Development and life course

- DV-01 · Grade A. Clauss JA, Blackford JU. *Behavioral inhibition and risk for developing social anxiety disorder: a meta-analytic study*. J Am Acad Child Adolesc Psychiatry. 2012;51:1066-1075.e1. DOI [10.1016/j.jaac.2012.08.002](https://doi.org/10.1016/j.jaac.2012.08.002); PMID [23021481](https://pubmed.ncbi.nlm.nih.gov/23021481/); PMCID [PMC3611590](https://pubmed.ncbi.nlm.nih.gov/PMC3611590/)
- DV-02 · Grade A. Sandstrom A et al. *Prospective association between childhood behavioral inhibition and anxiety: a meta-analysis*. J Abnorm Child Psychol. 2020;48:57-66. DOI [10.1007/s10802-019-00588-5](https://doi.org/10.1007/s10802-019-00588-5); PMID [31642030](https://pubmed.ncbi.nlm.nih.gov/31642030/)
- DV-03 · Grade B. Chronis-Tuscano A et al. *Stable early maternal report of behavioral inhibition predicts lifetime social anxiety disorder in adolescence*. J Am Acad Child Adolesc Psychiatry. 2009;48:928-935. DOI [10.1097/CHI.0b013e3181ae09df](https://doi.org/10.1097/CHI.0b013e3181ae09df); PMID [19625982](https://pubmed.ncbi.nlm.nih.gov/19625982/)
- DV-04 · Grade A. McLeod BD et al. *Examining the association between parenting and childhood anxiety: a meta-analysis*. Clin Psychol Rev. 2007;27:155-172. DOI [10.1016/j.cpr.2006.09.002](https://doi.org/10.1016/j.cpr.2006.09.002); PMID [17112647](https://pubmed.ncbi.nlm.nih.gov/17112647/)
- DV-05 · Grade A. Yap MBH et al. *Parental factors associated with depression and anxiety in young people: a systematic review and meta-analysis*. J Affect Disord. 2014;156:8-23. DOI [10.1016/j.jad.2013.11.007](https://doi.org/10.1016/j.jad.2013.11.007); PMID [24308895](https://pubmed.ncbi.nlm.nih.gov/24308895/)
- DV-06 · Grade B. Lebowitz ER et al. *Family accommodation in pediatric anxiety disorders*. Depress Anxiety. 2013;30:47-54. DOI [10.1002/da.21998](https://doi.org/10.1002/da.21998); PMID [22965863](https://pubmed.ncbi.nlm.nih.gov/22965863/); PMCID [PMC3932435](https://pubmed.ncbi.nlm.nih.gov/PMC3932435/)
- DV-07 · Grade A. Lebowitz ER et al. *Parent-based treatment as efficacious as CBT for childhood anxiety: a randomized noninferiority study of SPACE*. J Am Acad Child Adolesc Psychiatry. 2020;59:362-372. DOI [10.1016/j.jaac.2019.02.014](https://doi.org/10.1016/j.jaac.2019.02.014); PMID [30851397](https://pubmed.ncbi.nlm.nih.gov/30851397/); PMCID [PMC6732048](https://pubmed.ncbi.nlm.nih.gov/PMC6732048/)
- DV-08 · Grade A. Walkup JT et al. *Cognitive behavioral therapy, sertraline, or a combination in childhood anxiety*. N Engl J Med. 2008;359:2753-2766. DOI [10.1056/NEJMoa0804633](https://doi.org/10.1056/NEJMoa0804633); PMID [18974308](https://pubmed.ncbi.nlm.nih.gov/18974308/); PMCID [PMC2702984](https://pubmed.ncbi.nlm.nih.gov/PMC2702984/); NCT [NCT00052078](https://clinicaltrials.gov/ct2/show/study/NCT00052078)
- DV-09 · Grade A. James AC et al. *Cognitive behavioural therapy for anxiety disorders in children and adolescents*. Cochrane Database Syst Rev. 2020;11:CD013162. DOI [10.1002/14651858.CD013162.pub2](https://doi.org/10.1002/14651858.CD013162.pub2); PMID [33196111](https://pubmed.ncbi.nlm.nih.gov/33196111/); PMCID [PMC8092480](https://pubmed.ncbi.nlm.nih.gov/PMC8092480/)
- DV-10 · Grade B. Ginsburg GS et al. *Naturalistic follow-up of youths treated for pediatric anxiety disorders*. JAMA Psychiatry. 2014;71:310-318. DOI [10.1001/jamapsychiatry.2013.4186](https://doi.org/10.1001/jamapsychiatry.2013.4186); PMID [24477837](https://pubmed.ncbi.nlm.nih.gov/24477837/); PMCID [PMC3969570](https://pubmed.ncbi.nlm.nih.gov/PMC3969570/)

- DV-11 · Grade B. Pine DS et al. *The risk for early-adulthood anxiety and depressive disorders in adolescents with anxiety and depressive disorders*. Arch Gen Psychiatry. 1998;55:56-64. DOI [10.1001/archpsyc.55.1.56](https://doi.org/10.1001/archpsyc.55.1.56); PMID [9435761](https://pubmed.ncbi.nlm.nih.gov/9435761/)
- DV-12 · Grade A. Solmi M et al. *Age at onset of mental disorders worldwide: large-scale meta-analysis of 192 epidemiological studies*. Mol Psychiatry. 2022;27:281-295. DOI [10.1038/s41380-021-01161-7](https://doi.org/10.1038/s41380-021-01161-7); PMID [34079068](https://pubmed.ncbi.nlm.nih.gov/34079068/)

Distributed neurobiology

- NB-01 · Grade B. Etkin A, Wager TD. *Functional neuroimaging of anxiety: a meta-analysis of emotional processing in PTSD, social anxiety disorder, and specific phobia*. Am J Psychiatry. 2007;164:1476-1488. DOI [10.1176/appi.ajp.2007.07030504](https://doi.org/10.1176/appi.ajp.2007.07030504); PMID [17898336](https://pubmed.ncbi.nlm.nih.gov/17898336/); PMCID [PMC3318959](https://pubmed.ncbi.nlm.nih.gov/pmc/PMC3318959/)
- NB-02 · Grade A. Chavanne AV, Robinson OJ. *The overlapping neurobiology of induced and pathological anxiety: a meta-analysis of functional neural activation*. Am J Psychiatry. 2021;178:156-164. DOI [10.1176/appi.ajp.2020.19111153](https://doi.org/10.1176/appi.ajp.2020.19111153); PMID [33054384](https://pubmed.ncbi.nlm.nih.gov/33054384/); PMCID [PMC7116679](https://pubmed.ncbi.nlm.nih.gov/pmc/PMC7116679/)
- NB-03 · Grade A. Fullana MA et al. *Neural signatures of human fear conditioning: an updated and extended meta-analysis of fMRI studies*. Mol Psychiatry. 2016;21:500-508. DOI [10.1038/mp.2015.88](https://doi.org/10.1038/mp.2015.88); PMID [26122585](https://pubmed.ncbi.nlm.nih.gov/26122585/)
- NB-04 · Grade A. Fullana MA et al. *Fear extinction in the human brain: a meta-analysis of fMRI studies in healthy participants*. Neurosci Biobehav Rev. 2018;88:16-25. DOI [10.1016/j.neubiorev.2018.03.002](https://doi.org/10.1016/j.neubiorev.2018.03.002); PMID [29530516](https://pubmed.ncbi.nlm.nih.gov/29530516/)
- NB-05 · Grade B. Mobbs D et al. *When fear is near: threat imminence elicits prefrontal-periaqueductal gray shifts in humans*. Science. 2007;317:1079-1083. DOI [10.1126/science.1144298](https://doi.org/10.1126/science.1144298); PMID [17717184](https://pubmed.ncbi.nlm.nih.gov/17717184/); PMCID [PMC2648508](https://pubmed.ncbi.nlm.nih.gov/pmc/PMC2648508/)
- NB-06 · Grade C. Feinstein JS et al. 2013. DOI [10.1038/nn.3323](https://doi.org/10.1038/nn.3323); PMID [23377128](https://pubmed.ncbi.nlm.nih.gov/23377128/)
- NB-07 · Grade B. Hur J et al. *Anxiety and the neurobiology of temporally uncertain threat anticipation*. J Neurosci. 2020;40:7949-7964. DOI [10.1523/JNEUROSCI.0704-20.2020](https://doi.org/10.1523/JNEUROSCI.0704-20.2020); PMID [32958570](https://pubmed.ncbi.nlm.nih.gov/32958570/); PMCID [PMC7548695](https://pubmed.ncbi.nlm.nih.gov/pmc/PMC7548695/)
- NB-08 · Grade A. Zorn JV et al. *Cortisol stress reactivity across psychiatric disorders: a systematic review and meta-analysis*. Psychoneuroendocrinology. 2017;77:25-36. DOI [10.1016/j.psyneuen.2016.11.036](https://doi.org/10.1016/j.psyneuen.2016.11.036); PMID [28012291](https://pubmed.ncbi.nlm.nih.gov/28012291/)
- NB-09 · Grade A. Chalmers JA et al. *Anxiety disorders are associated with reduced heart rate variability: a meta-analysis*. Front Psychiatry. 2014;5:80. DOI [10.3389/fpsy.2014.00080](https://doi.org/10.3389/fpsy.2014.00080); PMID [25071612](https://pubmed.ncbi.nlm.nih.gov/25071612/); PMCID [PMC4092363](https://pubmed.ncbi.nlm.nih.gov/pmc/PMC4092363/)
- NB-10 · Grade B. Janiri D et al. *Shared neural phenotypes for mood and anxiety disorders: a meta-analysis of task-related brain activation studies*. JAMA Psychiatry. 2020;77:172-179. DOI [10.1001/jamapsychiatry.2019.3351](https://doi.org/10.1001/jamapsychiatry.2019.3351); PMID [31664439](https://pubmed.ncbi.nlm.nih.gov/31664439/)

Differential diagnosis and medical safety

- DD-01 · Grade A. Siegmann EM et al. *Association of depression and anxiety disorders with autoimmune thyroiditis: a systematic review and meta-analysis*. JAMA Psychiatry. 2018;75:577-584. DOI [10.1001/jamapsychiatry.2018.0190](https://doi.org/10.1001/jamapsychiatry.2018.0190); PMID [29800939](https://pubmed.ncbi.nlm.nih.gov/29800939/); correction DOI [10.1001/jamapsychiatry.2019.1643](https://doi.org/10.1001/jamapsychiatry.2019.1643)
- DD-02 · Grade A. Klevebrant L, Frick A. *Effects of caffeine on anxiety and panic attacks in patients with panic disorder: a systematic review and meta-analysis*. Gen Hosp

- Psychiatry. 2022;74:22-31. DOI [10.1016/j.genhosppsy.2021.11.005](https://doi.org/10.1016/j.genhosppsy.2021.11.005); PMID [34871964](https://pubmed.ncbi.nlm.nih.gov/34871964/)
- DD-03 · Grade A. Hindley G et al. *Psychiatric symptoms caused by cannabis constituents: a systematic review and meta-analysis*. Lancet Psychiatry. 2020;7:344-353. DOI [10.1016/S2215-0366\(20\)30074-2](https://doi.org/10.1016/S2215-0366(20)30074-2); PMID [32197092](https://pubmed.ncbi.nlm.nih.gov/32197092/); PMCID [PMC7738353](https://pubmed.ncbi.nlm.nih.gov/PMC7738353/)
- DD-04 · Grade C. Raj V et al. *Psychiatric profile and attention deficits in postural tachycardia syndrome*. J Neurol Neurosurg Psychiatry. 2009;80:339-344. DOI [10.1136/jnnp.2008.144360](https://doi.org/10.1136/jnnp.2008.144360); PMID [18977825](https://pubmed.ncbi.nlm.nih.gov/18977825/); PMCID [PMC2758320](https://pubmed.ncbi.nlm.nih.gov/PMC2758320/)
- DD-05 · Grade A. Tully PJ et al. *Panic disorder and incident coronary heart disease: a systematic review and meta-analysis*. Psychol Med. 2015. DOI [10.1017/S0033291715000963](https://doi.org/10.1017/S0033291715000963)
- DD-06 · Grade A. Tural U, Iosifescu DV. 2021. DOI [10.1016/j.jpsychires.2020.11.032](https://doi.org/10.1016/j.jpsychires.2020.11.032); PMID [33250190](https://pubmed.ncbi.nlm.nih.gov/33250190/)
- DD-07 · Grade B. Koning ASCAM et al. *Neuropsychiatric adverse effects of synthetic glucocorticoids: a systematic review and meta-analysis*. J Clin Endocrinol Metab. 2024;109:e1442-e1451. DOI [10.1210/clinem/dgad701](https://doi.org/10.1210/clinem/dgad701); PMID [38038629](https://pubmed.ncbi.nlm.nih.gov/38038629/); PMCID [PMC11099480](https://pubmed.ncbi.nlm.nih.gov/PMC11099480/)

Treatment, outcomes, and treatment resistance

- TX-01 · Grade A. Carpenter JK et al. *Cognitive behavioral therapy for anxiety and related disorders: a meta-analysis of randomized placebo-controlled trials*. Depress Anxiety. 2018;35:502-514. DOI [10.1002/da.22728](https://doi.org/10.1002/da.22728); PMID [29451967](https://pubmed.ncbi.nlm.nih.gov/29451967/); PMCID [PMC5992015](https://pubmed.ncbi.nlm.nih.gov/PMC5992015/)
- TX-02 · Grade A. van Dis EAM et al. *Long-term outcomes of cognitive behavioral therapy for anxiety-related disorders: a systematic review and meta-analysis*. JAMA Psychiatry. 2020;77:265-273. DOI [10.1001/jamapsychiatry.2019.3986](https://doi.org/10.1001/jamapsychiatry.2019.3986); PMID [31758858](https://pubmed.ncbi.nlm.nih.gov/31758858/)
- TX-03 · Grade A. Springer KS et al. *Remission in CBT for adult anxiety disorders: a meta-analysis*. Clin Psychol Rev. 2018;61:1-8. DOI [10.1016/j.cpr.2018.03.002](https://doi.org/10.1016/j.cpr.2018.03.002)
- TX-04 · Grade A. Papola D et al. *Psychotherapies for generalized anxiety disorder in adults: a systematic review and network meta-analysis*. JAMA Psychiatry. 2024;81:250-259. DOI [10.1001/jamapsychiatry.2023.3971](https://doi.org/10.1001/jamapsychiatry.2023.3971); PMID [37851421](https://pubmed.ncbi.nlm.nih.gov/37851421/); PMCID [PMC10585589](https://pubmed.ncbi.nlm.nih.gov/PMC10585589/)
- TX-05 · Grade A. Mayo-Wilson E et al. *Psychological and pharmacological interventions for social anxiety disorder in adults: a systematic review and network meta-analysis*. Lancet Psychiatry. 2014;1:368-376. DOI [10.1016/S2215-0366\(14\)70329-3](https://doi.org/10.1016/S2215-0366(14)70329-3); PMID [26361000](https://pubmed.ncbi.nlm.nih.gov/26361000/)
- TX-06 · Grade A. Pompoli A et al. *Psychological therapies for panic disorder with or without agoraphobia in adults: a network meta-analysis*. Cochrane Database Syst Rev. 2016;CD011004. PMID [27071857](https://pubmed.ncbi.nlm.nih.gov/27071857/)
- TX-07 · Grade B. Pompoli A et al. 2018. DOI [10.1017/S0033291717003919](https://doi.org/10.1017/S0033291717003919); PMID [29368665](https://pubmed.ncbi.nlm.nih.gov/29368665/)
- TX-08 · Grade A. Chawla N et al. *Which medications are best for the treatment of panic disorder with or without agoraphobia in adults? A network meta-analysis*. BMJ. 2022;376:e066084. DOI [10.1136/bmj-2021-066084](https://doi.org/10.1136/bmj-2021-066084); PMID [35045991](https://pubmed.ncbi.nlm.nih.gov/35045991/)

- TX-09 · Grade A. Snee A et al. *Pharmacological treatments for generalised anxiety disorder: a systematic review and network meta-analysis*. Lancet. 2019;393:768-777. DOI [10.1016/S0140-6736\(18\)31793-8](https://doi.org/10.1016/S0140-6736(18)31793-8); PMID [30712879](https://pubmed.ncbi.nlm.nih.gov/30712879/)
- TX-10 · Grade A. Guaiana G et al. *Pharmacological treatments in panic disorder in adults: a network meta-analysis*. Cochrane Database Syst Rev. 2023;CD012729. PMID [38014714](https://pubmed.ncbi.nlm.nih.gov/38014714/)
- TX-11 · Grade A. Batelaan NM et al. *Risk of relapse after antidepressant discontinuation in anxiety disorders, OCD, and PTSD: systematic review and meta-analysis of relapse prevention trials*. BMJ. 2017;358:j3927. DOI [10.1136/bmj.j3927](https://doi.org/10.1136/bmj.j3927); PMID [28903922](https://pubmed.ncbi.nlm.nih.gov/28903922/); PMCID [PMC5596392](https://pubmed.ncbi.nlm.nih.gov/PMC5596392/)
- TX-12 · Grade A. Normann N, Morina N. *The efficacy of metacognitive therapy: a systematic review and meta-analysis*. Front Psychol. 2018;9:2211. DOI [10.3389/fpsyg.2018.02211](https://doi.org/10.3389/fpsyg.2018.02211); PMID [30487770](https://pubmed.ncbi.nlm.nih.gov/30487770/); PMCID [PMC6246690](https://pubmed.ncbi.nlm.nih.gov/PMC6246690/)
- TX-13 · Grade A. Goldberg SB et al. *Mindfulness-based interventions for psychiatric disorders: a systematic review and meta-analysis*. Clin Psychol Rev. 2018;59:52-60. DOI [10.1016/j.cpr.2017.10.011](https://doi.org/10.1016/j.cpr.2017.10.011); PMID [29126747](https://pubmed.ncbi.nlm.nih.gov/29126747/)
- TX-14 · Grade A. Gower T et al. *Acceptance and commitment therapy versus bona fide psychotherapies: a systematic review and meta-analysis*. Clin Psychol Rev. 2026;128:102786. DOI [10.1016/j.cpr.2026.102786](https://doi.org/10.1016/j.cpr.2026.102786); PMID [42526063](https://pubmed.ncbi.nlm.nih.gov/42526063/)
- TX-15 · Grade A. Mataix-Cols D et al. *D-cycloserine augmentation of exposure-based CBT for anxiety, OCD, and PTSD: individual participant data meta-analysis*. JAMA Psychiatry. 2017;74:501-510. DOI [10.1001/jamapsychiatry.2016.3955](https://doi.org/10.1001/jamapsychiatry.2016.3955)
- TX-16 · Grade A. Bokma WA et al. *Aligning the many definitions of treatment resistance in anxiety disorders: a systematic review*. Depress Anxiety. 2019;36:801-812. DOI [10.1002/da.22895](https://doi.org/10.1002/da.22895); PMID [31231925](https://pubmed.ncbi.nlm.nih.gov/31231925/); PMCID [PMC6771798](https://pubmed.ncbi.nlm.nih.gov/PMC6771798/)
- TX-17 · Grade B. Domschke K et al. *International expert consensus on treatment-resistant anxiety disorders*. World Psychiatry. 2024. DOI [10.1002/wps.21177](https://doi.org/10.1002/wps.21177); PMID [38214637](https://pubmed.ncbi.nlm.nih.gov/38214637/)
- TX-18 · Grade B. Patterson B, van Ameringen M. *Augmentation strategies for treatment-resistant anxiety disorders: a systematic review and meta-analysis*. Depress Anxiety. 2016;33:728-736. DOI [10.1002/da.22525](https://doi.org/10.1002/da.22525)
- TX-19 · Grade B. Cirillo P et al. *Transcranial magnetic stimulation in anxiety and trauma-related disorders: a systematic review and meta-analysis*. Brain Behav. 2019;9:e01284. DOI [10.1002/brb3.1284](https://doi.org/10.1002/brb3.1284); PMID [31066227](https://pubmed.ncbi.nlm.nih.gov/31066227/); PMCID [PMC6576151](https://pubmed.ncbi.nlm.nih.gov/PMC6576151/)
- TX-20 · Grade B. Shinfuku M et al. *Effectiveness and safety of long-term benzodiazepine use in anxiety disorders*. Int Clin Psychopharmacol. 2019;34:211-221. DOI [10.1097/YIC.0000000000000276](https://doi.org/10.1097/YIC.0000000000000276); PMID [31274696](https://pubmed.ncbi.nlm.nih.gov/31274696/)
- TX-21 · Grade A. Schiele MA et al. *Integrative systematic review on pharmacological, psychotherapeutic, and neurostimulatory treatment options in treatment-resistant anxiety disorders*. Psychother Psychosom. 2026;95:214-247. DOI [10.1159/000547926](https://doi.org/10.1159/000547926); PMID [40946318](https://pubmed.ncbi.nlm.nih.gov/40946318/)