

LANDRY INDUSTRIES

A SCIENTIFIC FIELD GUIDE

The Teaching Signal

*The science of learning and evolving from
failure, mistake, and error*

Marie-Soleil Seshat Landry

FOUNDER & CEO, LANDRY INDUSTRIES

+ better than expected

- worse than expected



LANDRY INDUSTRIES

The Teaching Signal

*The science of learning and evolving from failure,
mistake, and error — across the brain, the person, and
the system*
written by

Marie-Soleil Seshat Landry

LANDRY INDUSTRIES · MONCTON, NEW BRUNSWICK

The Teaching Signal: The Science of Learning and Evolving from Failure, Mistake, and Error. First edition.

© 2026 Marie-Soleil Seshat Landry. Published by Landry Industries Inc., Moncton, New Brunswick, Canada. ORCID 0009-0008-5027-3337.

All rights reserved. This work synthesises peer-reviewed research and is intended as an educational reference. It is not a substitute for professional clinical, psychological, legal, or safety-engineering advice.

Verification note. Every empirical claim, named study, and quantitative figure in this book was checked against primary or authoritative secondary sources at the time of writing. Findings are attributed by author and year and listed in full in References. Where a body of evidence is contested, the disagreement is reported rather than smoothed over. Nothing here is cited that was not retrieved and read.

Set in **Fraunces** (display), **Source Serif 4** (text), and **Inter** (labels). Produced with an HTML→PDF typesetting pipeline.

A note on sensitivity: this book discusses error, failure, and — in the context of safety science — harm and mortality statistics. It also touches on self-criticism. If reflecting on personal failure becomes distressing, treat that as a signal to pause and, if needed, to reach for human support rather than to push on alone.

*His errors are the portals of
discovery.*

JAMES JOYCE · ULYSSES

Contents

	<i>Introduction — The Most Useful Thing That Can Happen to You</i>	7
PART ONE · THE NATURE OF ERROR		
1	The Inevitability of Error	13
2	A Field Taxonomy of Failure	17
3	The Brain That Runs on Error	22
PART TWO · THE SCIENCE OF LEARNING FROM ERROR		
4	Surprise Is Information	28
5	Desirable Difficulties	32
6	The Architecture of Practice	36
7	Mindset, Honestly	40
PART THREE · FROM PERSON TO SYSTEM		
8	The Mind's Cover-Ups	44
9	Safe to Fail	48
10	The Mindful Organization	52

PART FOUR · PRACTICE

11	The Error-Correction Protocol	58
12	Building a Learning Culture	62
13	Science, the Engine of Correction	66

CLOSING

	<i>Coda — The Discipline of Being Wrong Well</i>	70
	<i>Glossary of Key Terms</i>	72
	<i>References</i>	76
	<i>About the Author</i>	80

INTRODUCTION

The Most Useful Thing That Can Happen to You

Failure is a word that delivers a verdict. The science underneath it describes something far more useful: information about the gap between what you expected and what the world actually did.

There is a quiet violence in the word *failure*. It does not describe an event; it passes a sentence. To have failed is to have been measured and found wanting — a statement about who you are, not about what happened. Held that way, a mistake is something to hide, explain away, or survive. Almost everything that makes failure useful is lost in the flinch.

This book is built on a different premise, and it is not a motivational one. It is a mechanical one. Strip away the verdict and a failure is simply a *discrepancy*: a measured difference between what a system predicted and what the world returned. A wrong answer, a botched launch, a near-miss on a hospital ward, a model that misclassifies — each is a packet of information about where a model of the world is incomplete. The discrepancy is

not the enemy of progress. It is the raw material of it. The only real waste is a discrepancy that no one reads.

The claim of the chapters that follow is that learning from error is not a virtue you summon on a good day. It is a *skill*, and it rests on machinery we can describe with some precision — machinery that turns out to be remarkably similar at three very different scales.

At the scale of a neuron, a person, and an organization, learning runs the same loop: predict, act, compare, update.

Inside a single brain, learning is literally error-driven. When an outcome is better or worse than expected, midbrain dopamine neurons fire a burst or fall briefly silent, broadcasting a quantitative *teaching signal* proportional to how wrong the prediction was. Within about fifty milliseconds of a mistake, a structure called the anterior cingulate cortex throws up an electrical flag — the error-related negativity — and people whose brains flag errors more strongly tend to learn from them faster. The nervous system did not evolve to avoid being wrong. It evolved to *metabolize* being wrong into a better prediction.

At the scale of a person, the same logic shows up in the learning sciences. A generation of careful experiments has found that the conditions which make learning feel smooth — rereading, massed cramming, being handed the answer — produce shallow, brittle knowledge, while the conditions that feel like struggle — testing yourself, spacing practice, and even *guessing wrong before being corrected* — produce knowledge that lasts. Er-

rors, followed by feedback, are not detours around learning. Under the right conditions they *are* the learning.

And at the scale of an organization, the highest-performing institutions in aviation, medicine, firefighting, and software have converged on the same insight from the opposite direction: the way to get fewer catastrophic failures is to get *more* small ones into the open. They treat near-misses as gifts, hunt for weak signals, and — crucially — separate the act of explaining a failure from the act of blaming a person, because blame is the most reliable way ever invented to make information disappear.

Good failure, bad failure

None of this is a license to celebrate failure indiscriminately. The culture that says “fail fast, fail often” is exactly as unhelpful as the culture that says failure is unacceptable; both refuse to distinguish the failures worth having from the ones worth preventing. The organizational scientist Amy Edmondson, whose work threads through this book, draws the line cleanly: a *basic* failure is a preventable lapse in familiar territory, a *complex* failure is a systemic breakdown with many causes, and an *intelligent* failure is the honourable result of a well-reasoned venture into the unknown (Edmondson, 2023). The first two are to be engineered out. The third is the price of discovery, and a person or organization that has stopped producing intelligent failures has usually stopped reaching.

So learning from error is not the same as forgiving all error, and it is certainly not the same as inviting harm. The discipline is in the distinctions: knowing which failures to prevent and

which to provoke, and building the conditions under which the unavoidable ones actually teach.

THE NECESSARY CONDITION

One finding recurs at every scale and is worth holding onto from the start: **error teaches only when it is paired with feedback and only when it is recoverable.** A wrong guess with no correction entrenches the error. A mistake in a system that punishes the messenger produces concealment, not learning. The work of this book is largely the work of building that condition — feedback, safety, and recoverability — so that failure can do what it is uniquely able to do.

How this book is organized

The four parts move from the inside out. *Part One* establishes what error actually is — why it is structurally unavoidable, how to name its varieties, and how the brain is wired to learn from it. *Part Two* turns to the science of learning: how surprise carries information, why difficulty is often desirable, what really separates expert practice from mere repetition, and what an honest reading of the famous “mindset” research does and does not support. *Part Three* widens the lens to the obstacles — the biases that hide our own errors from us — and then to the systems that the best organizations build to surface error safely. *Part Four* is practical: a personal protocol for converting a failure into a lesson, a guide to building a team that learns, and a closing argument that science itself is nothing more, and nothing less, than humanity’s great institutional machine for being wrong well.

Throughout, boxes marked **Field Note** distill a principle into something you can act on, and figures render the key mechanisms visually. The aim is not to make you comfortable with failure. Comfort is beside the point. The aim is to make you *fluent* in it — able to read the signal in the noise, and to evolve.

PART ONE

I

The Nature of Error

Before failure can be useful, it has to be understood for what it is: not a moral event but a structural one. What it is, what kinds there are, and why the organ between your ears was built to run on it.

CHAPTER ONE

The Inevitability of Error

Mistakes are not occasional malfunctions in an otherwise correct system. For anything that acts under uncertainty, being wrong part of the time is a mathematical certainty — and the engine of every adaptation that has ever occurred.

Consider any system that has to act before it has complete information — which is to say, any system that acts at all. A chess player moving with the clock running, an immune cell deciding whether a protein is friend or foe, a company committing to a product two years before the market exists, a brain guessing what the next word in this sentence will be. None of them can wait for certainty, because certainty never arrives. Each must commit to a prediction on partial evidence, and a prediction made on partial evidence will sometimes be wrong. The wrongness is not a flaw in the system. It is a property of the situation.

We tend to experience our errors as personal — as evidence that we, specifically, are careless or slow. It is worth seeing clearly that error is built into the structure of the problem long before any individual shows up to commit one. There are simply more ways to be wrong than to be right. A combination

lock with four dials has ten thousand settings and one that opens; a sentence can be assembled in astronomically many ungrammatical orders and relatively few sensible ones; a protein can fold into countless useless shapes and a handful of functional ones. In almost every domain that matters, correct configurations are rare and incorrect ones are abundant. A system feeling its way toward the rare correct states will, by sheer arithmetic, spend much of its time in incorrect ones. Error is not the exception the universe makes for the incompetent. It is the default that competence has to climb out of.

The map is never the territory

There is a deeper reason error can never be driven fully to zero. Every system that acts intelligently does so by carrying a *model* — an internal map of how the world works — and using it to predict what will happen. But a map that captured every detail of the territory would be the size of the territory and just as useless. To be usable, a model must compress, simplify, and leave things out, and whatever it leaves out is a place where reality can surprise it. The semanticist Alfred Korzybski put the principle in a single line in 1933: the map is not the territory. The statistician George Box put the working consequence even more bluntly.

All models are wrong, but some are useful.

Box meant this not as cynicism but as the operating manual for anyone who builds models of anything — which, again, is everyone who thinks. The question is never whether your model is

wrong; it is wrong, guaranteed, in some respect you have not yet discovered. The only question is whether it is useful, and whether you have a way of finding the respects in which it is wrong so that you can make it more useful. Every error is a report from exactly that frontier — a notification that the territory has features your map does not. Read that way, a mistake stops being an indictment and becomes a survey instrument. It tells you precisely where your model and the world diverge, which is the only place improvement was ever going to come from.

Evolution is an edited mistake

If this sounds like a fragile way to build anything, consider that it is the only way nature has ever built anything at all. Biological evolution is, mechanically, a process of error and feedback: genetic variation is generated partly by *copying mistakes* in replication, and natural selection is the feedback that keeps the rare useful errors and discards the rest. Every organism alive, including the one reading this, is the cumulative product of billions of years of edited mistakes. The immune system runs the same algorithm in miniature, generating enormous random diversity in antibodies and selecting the few that happen to fit a threat. Trial and error is not a primitive method that better methods have replaced; across biology it is the *only* general-purpose method for finding solutions in a space too large and too unknown to reason through in advance. Variation proposes, the world disposes, and what survives is knowledge.

This is why the chapters ahead keep returning to a single loop — predict, act, compare, update — at the scale of neurons, individuals, and institutions. It is the same loop evolution runs

across generations, and it is the only one that demonstrably works when the territory is genuinely unknown. The systems we admire for “not making mistakes” are almost never error-free. They are *error-processing*: arranged so that mistakes are small, surfaced quickly, and fed back into a better model before they compound.

FROM VERDICT TO VARIABLE

The single most useful shift this book asks for is to stop treating a mistake as a **verdict** on your worth and start treating it as a **variable** in a model you are improving. A verdict is final and is about you. A variable is informative and is about the world. The same event — the failed pitch, the rejected paper, the bug in production — can be either, depending entirely on what you do with it next. Nothing about the event decides which it is. You do.

To accept the inevitability of error is not to lower one’s standards; it is to aim them correctly. The goal of a serious person or a serious organization is not to never be wrong, which is impossible, but to be wrong in *survivable, informative, ever-smaller* ways — and to extract the maximum signal from each. The rest of this book is about how. We begin where any diagnosis begins: by learning to tell one kind of failure from another.

CHAPTER TWO

A Field Taxonomy of Failure

An emergency physician never treats “illness.” They diagnose. Treating every failure as one undifferentiated thing is the surest way to learn nothing from any of them — because each kind carries a different lesson and demands a different response.

The word *failure* is a poor diagnostic instrument. It lumps together a typo and a plane crash, a missed sales target and a falsified data set, a well-designed experiment that didn't pan out and an act of sabotage. These have almost nothing in common except that the outcome was undesired, and they call for responses as different as an apology, an engineering change, a celebration, and a dismissal. Before reacting to a failure, the disciplined move is the diagnostic one: *which kind is this?* The answer determines what, if anything, there is to learn.

Slips, lapses, and mistakes

The psychologist James Reason spent decades studying how competent people fail, and his foundational distinction, set out in *Human Error* (1990), is between failures of *execution* and failures of *intention*. A *slip* or *lapse* occurs when the plan was correct but the action went astray: you intended to stop at the phar-

macy but drove your usual route home on autopilot; you knew the password but your fingers typed last year's. The intention was sound; attention or memory failed to carry it out. A *mistake*, in Reason's precise sense, is the opposite and the more dangerous: the action was carried out exactly as planned, but *the plan itself was wrong*. You followed the recipe perfectly; the recipe was bad.

The distinction matters because the remedies are opposites. Slips and lapses are failures of a basically correct system under load, and they are reduced by better design — checklists, constraints, reducing distraction and fatigue, building the right action into the path of least resistance. Mistakes are failures of knowledge or reasoning, and no amount of vigilance fixes them; they are reduced only by better models, better information, and better feedback. Punishing a person for a slip is unjust and useless, because the next tired person on the same badly designed system will slip in the same place. Treating a knowledge-based mistake as mere carelessness is worse, because it leaves the faulty model in place to fail again.

Basic, complex, and intelligent


Where Reason classified failures by their cognitive mechanism, the Harvard organizational scientist Amy Edmondson classifies them by their relationship to *learning* — a framework she developed across decades of research and crystallized in *Right Kind of Wrong* (2023). She names three archetypes. A **basic failure** is a preventable error in familiar territory, where a known correct method existed and was not followed — the kind that produces no new knowledge and is simply to be prevented. A

complex failure arises in a complicated system from several causes lining up at once: no single blunder, but a cascade, the sort that produces hospital adverse events and industrial accidents. And an **intelligent failure** is the honourable outcome of a thoughtful foray into genuinely new territory — a well-designed experiment, a reasonable bet, that did not work.

Edmondson's point is that our two reigning cultural attitudes both fail because they refuse this distinction. "Failure is unacceptable" suppresses the intelligent failures that are the only path into the unknown. "Fail fast, fail often" romanticizes failure indiscriminately and excuses the basic ones that diligence should have prevented. The mature stance is neither: *prevent* the basic failures, *untangle* the complex ones, and *provoke* more intelligent ones. An organization that has stopped producing intelligent failures has not achieved excellence; it has stopped reaching.

The spectrum of blame

Underneath these categories runs a question every team handles badly: *was someone to blame?* Edmondson's answer, drawn from her fieldwork, is that genuinely blameworthy failure is far rarer than our instincts insist. She arranges the reasons a failure happens along a spectrum (Edmondson, 2011), from the clearly culpable to the clearly praiseworthy.



Truly culpable failure — the left end — is uncommon. Most failure sits in the middle (systems and circumstance) or, at the right, is the unavoidable cost of doing something genuinely new.

Figure 2.1. A spectrum of reasons for failure, after Edmondson (2011). The instinct to find a culprit assumes the failure belongs on the far left; the evidence is that it rarely does.

Read the spectrum from left to right and a striking fact emerges: only the leftmost reasons — deliberate sabotage, recklessness — are genuinely blameworthy, and these are rare. The large middle band is occupied by inattention, gaps in ability, and inadequate processes — failures of the *system* far more than of the will. And the right end, where someone tried something reasonable that happened not to work, is not blameworthy at all; it is the texture of progress. Yet when surveyed about their own organizations, managers consistently estimate that only a small fraction of failures are truly blameworthy — while admitting they treat a large majority *as if* they were. That gap, between how few failures deserve blame and how many receive it, is where most organizational learning goes to die.

CLASSIFY BEFORE YOU REACT

When a failure lands, the untrained response is to assign feeling (frustration, shame, blame) and move on. The trained response is to assign a **category** first. Ask, in order: Was the plan wrong, or only its execution? Is this a basic failure that a known method would have prevented, a complex one with many causes, or an intelligent one at the edge of the known? And honestly — where on the blame spectrum does it actually sit? Only after the failure is classified does its lesson become legible. Misclassify it, and you will “fix” the wrong thing: drilling vigilance into a knowledge problem, or blaming a person for a system.

A taxonomy is not bureaucratic tidiness; it is the difference between learning and merely suffering. The same wrong outcome can be a prompt to redesign a process, to revise a belief, to celebrate a smart bet, or — occasionally — to part ways with someone. Telling them apart is the first skill of failing well. The next chapter goes a layer deeper, into the organ that performs this metabolism of error every waking second, usually without telling you.

CHAPTER THREE

The Brain That Runs on Error

Learning is not something the brain does in spite of mistakes. At the level of cells and electrical signals, learning is the correction of mistakes — and the same mathematics turns out to drive both nervous systems and the machines we build to learn.

In a series of experiments that reshaped neuroscience, the physiologist Wolfram Schultz recorded from single dopamine neurons in the midbrain while an animal learned that a cue predicted a reward. He expected the neurons to fire when the reward arrived. Instead, he found something stranger and far more revealing. Early in learning, before the animal knew what the cue meant, the neurons fired when the reward appeared — an *unexpected* good thing. But once the animal had learned that the cue reliably predicted the reward, the neurons stopped firing at the reward and fired at the *cue* instead. And if the predicted reward was then withheld, the neurons fell briefly *silent* at the exact moment it should have come, dipping below their baseline rate (Schultz, Dayan, & Montague, 1997).

The neurons were not signalling reward. They were signalling *surprise about* reward — the difference between what was predicted and what occurred. Neuroscientists call this quantity

the **reward prediction error**, and dopamine is its currency: a burst above baseline when the world is better than expected (a positive error), no change when the world is exactly as predicted (no error, nothing to learn), and a dip below baseline when the world is worse than expected (a negative error). The size of the deflection encodes the size of the surprise. It is, in the most literal sense the brain has, a *teaching signal*.




Figure 3.1. The dopamine teaching signal, after Schultz, Dayan & Montague (1997).

The cells report the error in prediction, not the reward itself. A fully expected outcome — however good — produces no signal, because there is nothing to learn. Surprise is what the brain is built to register.

This is not a metaphor borrowed from computing; the causality runs the other way. The mathematics Schultz’s neurons turned out to implement is the same *temporal-difference* rule that computer scientists Richard Sutton and Andrew Barto had developed for machine learning — the principle that an agent improves by adjusting its predictions in proportion to its prediction errors. When today’s artificial systems learn by reinforcement, they are running a refined version of the algorithm a monkey’s midbrain was already running in the 1990s. “Learning

from error” is not an inspirational slogan that happens to apply to brains and machines alike. It is the literal shared substrate of both.

The fifty-millisecond flag

Prediction error is how the brain learns from outcomes. A second system handles a more immediate job: catching errors in our own actions, often before we are consciously aware of them. In 1993, a team led by William Gehring discovered a sharp electrical signature that appears over the front of the scalp within roughly fifty milliseconds of a person making a mistake in a fast task — faster than conscious recognition. They named it the **error-related negativity**, and it is generated by the anterior cingulate cortex, a region sitting at the junction of thought and motivation (Gehring et al., 1993). Before you know you have erred, your brain already does, and is already mobilizing to correct.

What ties the two systems together is a striking theory proposed by Clay Holroyd and Michael Coles (2002): the error-related negativity is what it looks like when a *negative* dopamine prediction-error signal — the dip that says “worse than expected” — reaches the anterior cingulate and is put to use adjusting behaviour. The flag that says “you just made a mistake” and the teaching signal that says “the world was worse than you predicted” are, on this account, two faces of one error-correction machine.

Who learns from their mistakes

If the error-related negativity is the brain registering its own mistakes, a natural question follows: do people whose brains flag errors more strongly actually *learn* from them more? The evidence says yes. Michael Frank and colleagues found that individuals with larger error-related negativities were better at learning to *avoid* choices that had led to bad outcomes — and that people differ reliably in whether they learn more from their errors or from their successes (Frank, Woroch, & Curran, 2005). The capacity to be instructed by a mistake is not uniform; it is, in part, a measurable property of a nervous system — and, as later chapters argue, a trainable one. The brain supplies the signal. Whether that signal becomes learning depends on what happens next: on attention, on feedback, and on whether it is safe to have erred at all.

PREDICT, THEN CHECK

There is an immediate, practical consequence of how this machinery works. A prediction error can only be computed if a prediction was actually made. When you read an answer without first trying to produce it, or watch an expert without first guessing what they will do, you hand your brain the outcome with no prediction to compare it against — and generate no teaching signal. The fix is almost absurdly simple and is one of the best-supported moves in all of learning science: **commit to a prediction before you check**. Guess the answer, forecast the result, call the outcome out loud. Being wrong then is not a failure of the exercise. It is the exercise.

We have established the terrain: error is structurally unavoidable, it comes in distinct kinds that demand distinct responses, and the brain is wired from the cellular level up to learn by computing how wrong it was. That is the foundation. The next part asks the question that foundation makes urgent — given that we are built to learn from error, how do we do it deliberately, and well?

PART TWO

II

Learning on Purpose

If the brain is built to learn from error, the next question is practical: how do we put that machinery to deliberate use? The science of how difficulty, surprise, and feedback turn mistakes into durable skill — and what the evidence really says about the methods we are sold.

CHAPTER FOUR

Surprise Is Information

A prediction that comes true teaches you nothing you did not already know. It is the prediction that fails — the moment the world refuses to match your model — that carries every bit of information worth having.

There is a strange idea at the centre of modern brain science, and it inverts the common-sense picture of perception. We tend to imagine the senses as cameras: light and sound stream in, and the brain assembles them into a picture of the world. The evidence increasingly says it works the other way around. The brain is constantly generating predictions about what it is about to sense, and what flows *up* from the senses is not the raw world but the *error* — the mismatch between what was predicted and what arrived. On this account, summarized by the philosopher Andy Clark, brains are “essentially prediction machines,” bundles of cells that perceive and act by continually trying to match incoming signals against top-down expectations and minimizing the difference (Clark, 2013).

The idea is older than its current vocabulary. In the nineteenth century Hermann von Helmholtz proposed that perception is a form of *unconscious inference*: the brain does not receive the world so much as guess it, building the best hypothesis it can and testing that hypothesis against the incoming

data. The modern framework — variously called predictive processing or the Bayesian brain — gives the guess a mathematical shape and the testing a mechanism. A hierarchy of predictions runs top-down; sensory evidence runs bottom-up; and the only thing that needs to travel upward is the part the prediction got wrong. Where the model is right, there is nothing to send. Surprise is the signal.

Why a confirmed expectation is silent

This has a consequence that sounds like a paradox but is simply arithmetic. The more confidently and correctly you predict something, the less you learn when it happens. A confirmed expectation contains no new information, because you already had it — that is what it means for the expectation to have been correct. The engineer Claude Shannon made this precise in founding information theory: the information carried by an event is inversely related to its probability. A certainty carries no information; a shock carries a great deal. The brain's teaching signal and the mathematician's definition of information are pointing at the same thing from two directions. What is worth learning from is, by definition, what you did not see coming.

The events that teach you the most are precisely the ones you were most confident would go the other way.

Turn that around and it becomes one of the most useful working rules a learner can adopt. If information lives in the gap

between prediction and outcome, then the size of that gap is a measure of how much there is to learn — and the events that teach you the most are precisely the ones you were most confident would go the other way. The expert who is mildly surprised has found a small correction to make. The expert who is *astonished* has found something genuinely wrong with their model, which is the most valuable thing an inquiry can turn up. A surprising result is not an annoyance interrupting the work. On a real account of how learning works, it *is* the work, arriving.

The forecast you can be graded on

None of this machinery can run, though, on a prediction that was never made. The brain can only compute the gap between expectation and reality if an expectation was actually committed to in advance. This is why so much of what feels like learning produces so little of it: we read the explanation, nod along, and absorb the outcome without ever having staked a prediction it could correct. The fix is to make the prediction explicit and falsifiable — to put a number, a guess, or a forecast on the record *before* the answer arrives, so that being wrong produces a signal instead of a shrug.

WRITE THE NUMBER DOWN

Before you check the result of an experiment, open a file, or hear how the meeting went, commit to a specific prediction — ideally in writing, ideally with a confidence attached. “I think this returns about 40 milliseconds; 70% sure.” When the answer comes, the distance between your forecast and the truth is a measured teaching signal you can actually feel. Over time the practice does something subtler still: it calibrates your sense of your own certainty, teaching you not just facts but how much to trust your beliefs about them. A forecast that cannot be graded is a wish. A forecast you have written down is an instrument.

So surprise is the nutrient of learning, and an unhedged prediction is what lets the body absorb it. The natural next move is almost mischievous: if violated expectations are what build durable knowledge, then good learning should be designed to *generate* them — to make the learner struggle, hesitate, and get things wrong on purpose. That runs against every instinct that equates smooth performance with progress. It is also one of the best-established findings in the science of learning, and it is where we turn next.

CHAPTER FIVE

Desirable Difficulties

The conditions that make learning feel easy and fast are often the ones that make it vanish. The conditions that make it feel slow and effortful — even error-ridden — are frequently the ones that make it last.

For decades, the psychologists Robert Bjork and Elizabeth Bjork have studied a counterintuitive and well-replicated phenomenon they named **desirable difficulties**: obstacles introduced into the learning process that slow down acquisition and make performance look worse in the moment, yet produce deeper understanding and far more durable retention over time. The catch — the reason the effect is so easy to miss — is that the difficulty *feels* like it is hindering you exactly when it is helping you most. We are strongly inclined to judge learning by how fluent it feels right now, and on that gauge, desirable difficulties read as failure.

Four of these difficulties are particularly well supported. *Retrieval practice*: testing yourself on material, struggling to pull it from memory, builds far stronger retention than rereading it, even though rereading feels more productive. *Spacing*: distributing study across time, allowing partial forgetting between sessions, beats massing it together, even though massing produces a more satisfying sense of mastery in the moment. *Interleaving*: mixing different problem types together rather than practising

one kind in a block, which feels confusing and error-prone but teaches the crucial skill of recognizing *which* approach a new problem calls for. And *generation*: trying to produce an answer before being shown it — guessing, attempting, getting it wrong — rather than simply studying the correct response.

The error before the answer

That last one carries the central message of this book to its sharpest point. It is natural to assume that making errors while studying is at best wasted effort and at worst actively harmful — that you might burn the wrong answer into memory. The evidence runs the other way. In a series of experiments, Nate Kornell, Matthew Hays, and Robert Bjork showed that *unsuccessful* retrieval attempts — trying to answer a question you cannot yet answer, and failing — enhanced subsequent learning of the answer, compared with spending the same time simply studying it (Kornell, Hays, & Bjork, 2009). The failed guess was not wasted. It primed the mind to absorb the correction, much as a prediction must be made before its error can teach. The struggle to generate an answer, even a wrong one, prepares the ground that the right answer then takes root in.




Figure 5.1. The signature of a desirable difficulty: a crossover interaction. Easy study wins on the immediate test and loses on the delayed one; effortful, error-prone study shows the reverse. Judged in the moment — the only moment the learner usually has — the better long-term method looks like the worse one.

The one condition that makes it work

There is, however, a non-negotiable condition, and ignoring it turns this whole chapter into dangerous advice. Errors enhance learning *only when they are followed by corrective feedback*. A wrong guess that is never corrected does not prime anything; it simply rehearses the wrong answer, and rehearsal is how memories strengthen. The benefit of the failed attempt comes entirely from the loop being closed — from the error being met, promptly, by the truth. Generate, then check. Struggle, then resolve. An error left to stand is not a desirable difficulty. It is just a mistake being practised.

This also explains one of the most persistent illusions in all of self-directed learning, what the Bjorks call the *illusion of fluency*. Rereading a chapter, watching a lecture again, reviewing notes we have highlighted — these feel like learning because the ma-

terial grows easy to process. But ease of processing is not the same as durability of memory, and we routinely mistake the first for the second. The student who rereads feels more confident and performs worse than the student who self-tests and feels worse. We are, in effect, equipped with a faulty gauge that reads *fluency* when we need it to read *retention* — which is exactly why the methods that work feel wrong.

WHEN DIFFICULTY STOPS BEING DESIRABLE

The principle has a real boundary, and overshooting it does harm. A difficulty is only desirable if the learner has the background knowledge to meet it and eventually succeed. Pile on obstacles a person cannot yet overcome — material miles beyond their level, or conditions so stressful that anxiety swallows the working memory the task requires — and the difficulty becomes simply undesirable: it produces frustration and failure with no compensating gain. The art is calibration. The struggle has to be the kind the learner can win, eventually, with effort. Productive difficulty sits just past the edge of current ability, not over the horizon.

Desirable difficulties tell us how to structure the encounter with material so that errors do their work. But they say little about how a whole skill is built over years — how a novice becomes an expert, and what role mistakes and feedback play across thousands of hours of practice. That question has its own large and frequently distorted literature, and getting it right means separating what the science actually shows from the story that got told about it.

CHAPTER SIX

The Architecture of Practice

Repetition is not practice. What turns hours into skill is a particular structure — working at the edge of your ability, where errors are frequent, and correcting them with immediate feedback. And what the famous number really meant is not what you were told.

In 1993, the psychologist K. Anders Ericsson and his colleagues published a study of violinists at a Berlin music academy that would become one of the most cited — and most distorted — findings in the science of expertise (Ericsson, Krampe, & Tesch-Römer, 1993). Comparing students of different attainment, they found that the best performers had accumulated far more hours of a specific kind of activity Ericsson called **deliberate practice**. Crucially, deliberate practice is not playing the pieces you already know, and it is not logging hours. It is effortful, goal-directed work on the specific things you *cannot yet do*, undertaken at the edge of your current ability, guided continuously by feedback, with errors identified and corrected in tight loops.

Read carefully, Ericsson's account is a theory about error. The reason ordinary repetition stops improving us — the reason a person can drive for thirty years without becoming a better

driver — is that it operates entirely within the zone of what we have already mastered, where mistakes are rare and feedback is absent. Deliberate practice deliberately leaves that comfortable zone. It seeks out the difficulty where errors are frequent, treats each error as information about what to fix next, and closes the loop before moving on. The structure that builds expertise is, at bottom, the structure for generating and correcting mistakes at a sustainable rate.

The number that ate the science

You have almost certainly heard that mastery takes ten thousand hours. That figure comes not from Ericsson but from the journalist Malcolm Gladwell, who in his 2008 book *Outliers* distilled the violin study into a memorable “10,000-hour rule.” The simplification was potent and almost entirely misleading. Ericsson himself repeatedly objected that the specific number was essentially arbitrary — an average for one group in one domain, not a threshold — and, more importantly, that the rule erased the entire content of his theory. Ten thousand hours of unfocused repetition produces ten thousand hours of no improvement. What matters is not the quantity of practice but its *quality*: whether it is structured to confront error at the edge of ability with feedback. The popular version kept the hours and threw away the only part that did the work.

Ten thousand hours of unfocused repetition produces ten thousand hours of no improvement.

How much practice really explains

If the popular story overstated practice in one direction, careful evidence has since corrected it in the other. In a meta-analysis spanning every major domain in which deliberate practice had been studied, Brooke Macnamara, David Hambrick, and Frederick Oswald asked how much of the variation in performance practice actually accounts for. Their answer was sobering: deliberate practice explained about 26% of the variance in games, 21% in music, 18% in sports, just 4% in education, and less than 1% in the professions. Their conclusion was measured and worth quoting in spirit: deliberate practice is important, but not as important as had been argued (Macnamara, Hambrick, & Oswald, 2014).

What fills the remaining variance is contested — some mix of starting age, working-memory capacity, the quality rather than mere quantity of practice, opportunity, coaching, and factors not yet well measured. The point is not that practice does not matter; in every domain it mattered, and in skill-based pursuits like games and music it mattered a great deal. The point is that the honest figure is a fraction, not the whole, and a culture that tells people effort is everything sets them up to read the ordinary limits of practice as personal failure. Ericsson and his collaborators dispute aspects of this meta-analysis to this day — which is exactly as it should be. The disagreement is being conducted in the open, in the currency of evidence, and the field's estimate of how much practice explains is more accurate now because of it.

PRACTICE AT THE EDGE, WITH A MIRROR

To make practice deliberate, two ingredients are non-negotiable, and both are about error. First, work at the boundary of your ability — on the specific sub-skill that breaks down, not the whole performance you can already get through. If you are not making mistakes, the practice is too easy to be teaching you anything. Second, install a feedback loop fast and honest enough to catch those mistakes: a coach, a recording, a test, an automated check, a metric. Effort aimed at your weakest point, held up to a mirror that shows you the error immediately — that is the engine. Hours without it are just time.

The architecture of practice tells us how skill is built once a person is willing to confront their errors. But willingness is exactly the variable in question. Whether someone treats a mistake as information to act on or as a verdict to flee from depends on what they believe failure *means* — and that belief has become the subject of one of the most popular, and most oversold, ideas in modern psychology.

CHAPTER SEVEN

Mindset, Honestly

The belief that ability can grow does change how a person meets failure. It is also a smaller, more fragile, and more honest effect than the bestseller version promised — and telling the truth about it matters more than defending the slogan.

The psychologist Carol Dweck drew a distinction that has reshaped how a generation talks about learning. A person with a *fixed mindset* believes their abilities are essentially set — you are either smart at this or you are not — and therefore experiences a mistake as exposure, evidence of a limit they cannot move. A person with a *growth mindset* believes abilities can develop through effort and good strategy, and therefore experiences the same mistake as information, a signpost toward what to practise next. Stated at this level, the idea is both plausible and consonant with everything in this book: how you interpret an error determines whether you can learn from it.

The empirical anchor most often cited is a 1998 experiment by Claudia Mueller and Carol Dweck. Children praised for intelligence after success (“you must be smart”) later avoided challenging tasks, gave up more easily after failure, and even misreported their scores; children praised for effort (“you must have worked hard”) sought out harder problems and persisted (Mueller & Dweck, 1998). The lesson drawn was elegant: praise

the controllable process, not the fixed trait, and you teach a child that failure is about strategy rather than worth.

What the larger evidence shows

Then the idea became an industry, and the claims outran the data. When Victoria Sisk and colleagues subjected the field to two large meta-analyses, the results were deflating for anyone who had been promised transformation. Across 273 effect sizes and more than 365,000 students, the overall relationship between a person's mindset and their academic achievement was weak — mindset accounted for roughly one percent of the variation in achievement. And across the intervention studies, programs designed to instill a growth mindset produced only a small average effect on achievement (Sisk et al., 2018). The headline version — that teaching children a belief about malleable intelligence reliably and substantially raises their performance — was not supported.

But the meta-analyses did not say the idea was empty, and honesty cuts both ways. Some results supported specific, narrower tenets of the theory: students who were academically at risk, or from low-socioeconomic-status backgrounds, did appear to benefit from mindset interventions in a way their more advantaged peers did not. A large national experiment by David Yeager and colleagues, conducted after the meta-analyses, told a similar story — a brief online intervention produced a small but real improvement, concentrated among lower-achieving students and in schools whose norms made the message believable. The realistic picture is neither the bestseller nor its back-

lash. The effect is small, it is conditional, and it lands hardest for the students with the most room to gain.

THE DIFFERENCE BETWEEN ENCOURAGEMENT AND DENIAL

There is a degraded version of this idea that does harm, and it should be named. Telling a struggling person that they can achieve anything if they simply believe, while changing nothing about their instruction, feedback, or circumstances, is not a growth mindset; it is an alibi for not helping them. Worse, it relocates failure back inside the person — if effort and belief are everything, then to still be failing is to have not believed hard enough. The evidence-based core is narrower and kinder: abilities really can improve with the right effort and strategy, mistakes really are part of how that improvement happens, and the useful response to a setback is to change the approach, not to indict the self. Encouragement supplies a better method. Denial just supplies a brighter feeling on the way to the same wall.

What survives the scrutiny, then, is modest and durable: how a person frames failure genuinely shapes whether they engage with it or flee, and the framing can be taught, especially to those who most need it. That is worth keeping. It is also where the individual story reaches its limit. For all the focus on private beliefs and personal practice, most consequential errors are not made by lone individuals at all. They are made inside teams, institutions, and systems — and whether those errors are surfaced and corrected or buried and repeated depends far less on any one person's mindset than on the conditions the system creates around being wrong.

PART THREE

III

The Error in the System

Most consequential mistakes are not made by lone individuals but inside teams and institutions. Why our minds hide our errors from us, why blame guarantees they stay hidden, and how the safest organizations in the world turn failure into the raw material of reliability.

CHAPTER EIGHT

The Mind's Cover-Ups

Before an error can be corrected it has to be seen, and the mind is equipped with a remarkable set of devices for not seeing it. These biases are not random malfunctions. They share a purpose: to protect us from the discomfort of having been wrong.

Everything in the previous chapters assumed a learner willing to look squarely at their own mistakes. That assumption is generous. The human mind comes furnished with a suite of cognitive biases whose net effect is to hide our errors from us — to edit the record, reassign the blame, and screen out the evidence, all below the level of awareness and all in the service of a self that would rather not have erred. To learn from error, you first have to defeat the machinery built to conceal it.

Consider **hindsight bias**, the “knew-it-all-along” effect first demonstrated by Baruch Fischhoff in the 1970s. Once we know how something turned out, we systematically misremember having expected that outcome all along. The genuine uncertainty we felt beforehand — the very thing that made the situation hard — is overwritten by the knowledge of the result. The damage to learning is precise: if you believe you knew it all along, there is nothing to learn, because by your own revised account you made no error. Hindsight bias does not just distort memory; it quietly cancels the lesson.

Consider the **self-serving bias**, the tendency to attribute our successes to our own skill and our failures to bad luck or difficult circumstances. It feels like fairness from the inside and functions as an exemption: by routing every failure to an external cause, it ensures we never have to revise the thing we control. Its social cousin, the **fundamental attribution error**, completes the asymmetry — when others fail we blame their character, and when we fail we blame our situation. The same event becomes a flaw in them and a misfortune to us. This double standard is the cognitive engine of blame, and, as the next chapter shows, blame is precisely what keeps organizational errors buried.

Consider **confirmation bias**, our pull toward evidence that supports what we already believe and away from evidence that would refute it. It is the exact opposite of the disposition that learning requires. Where the scientific attitude actively hunts for the disconfirming case — the result that would prove the cherished hypothesis wrong — confirmation bias steers us to gather a comfortable pile of agreement and call it proof. A belief tested only against friendly evidence has not been tested at all.

The errors you never see

The most insidious of these distortions is the one that operates on the data itself. During the Second World War, the statistician Abraham Wald was asked by the U.S. military's Statistical Research Group to advise on where to add armour to bombers. The planes returning from missions showed dense patterns of bullet holes across the wings and fuselage and few on the engines, and the natural recommendation was to reinforce where the holes

clustered. Wald saw the fatal flaw. The data came only from planes that had *returned*. Planes hit in the engines were not in the sample because they had not come back. The undamaged areas on the survivors marked precisely the places that were lethal when struck — so the armour belonged where the holes were not (Wald, via Mangel & Samaniego, 1984).

The planes that could have taught the most important lesson were the ones that never came back to tell it.

This is **survivorship bias**, and it is everywhere learning draws its lessons only from successes. We study the habits of billionaire founders and conclude that dropping out of college builds companies, ignoring the unseen multitude who dropped out and built nothing. We reverse-engineer the routines of the people who made it and miss that the people who did the identical things and failed have been quietly removed from view. The failures contain the information — often the *decisive* information — and they are exactly what the world stops showing us. A learner who attends only to what survived is studying a map with the most important territory torn off.

These biases share a single, uncomfortable function: they defend the ego against the pain of being wrong, and they do it by disabling the error-correction machinery that the earlier chapters described. The first defence against them is simply knowing they are there — naming the cover-up as it happens. But self-awareness is a weak and tiring tool to run against forces this automatic. The durable correction is not internal but struc-

tural: building environments and routines that surface errors whether or not any individual is willing to see their own. That is a problem of design, and design begins with the single condition that makes everything else possible — whether it is safe to be wrong out loud.

CHAPTER NINE

Safe to Fail

An error that is hidden cannot be corrected. The single greatest determinant of whether a team learns from its mistakes is not the talent of its members but whether they believe it is safe to admit a mistake at all.

In the early 1990s, a young researcher named Amy Edmondson set out to test an intuitive hypothesis: that better hospital teams make fewer medication errors. Her data said the opposite. The best-rated nursing teams appeared to make *more* errors, not fewer. The result made no sense until she looked closer and found the explanation that would shape the next thirty years of her work. The better teams were not making more mistakes; they were *reporting* more, because their members felt safe enough to admit them. On the worse teams, errors were happening just as often and being quietly hidden. The reported error rate was not measuring incompetence. It was measuring candour.

Edmondson named the underlying condition **psychological safety**: a shared belief that the team is a safe place for interpersonal risk-taking — that you can ask a question, admit a mistake, or raise a concern without being humiliated or punished for it. Years later, when Google ran a two-year study of what makes its teams effective, code-named Project Aristotle, it analysed scores of variables across some 180 teams and found that the strongest

predictor of team performance was not the seniority, intelligence, or individual talent of the members. It was psychological safety — whether, in the words of the survey item that best captured it, making a mistake on the team would be held against you (Edmondson, 1999; Google re:Work, Project Aristotle).

Why blame guarantees concealment

The mechanism is not mysterious. When an error is met with blame, the rational response is to hide the next one. Concealment is not a character defect; it is what punishment teaches. And a hidden error is an error that cannot be analysed, cannot be corrected, and is therefore free to happen again — usually to someone else, usually worse. An organization that punishes the messenger does not reduce its error rate. It reduces its *knowledge* of its error rate, which is far more dangerous, because it now believes itself safe while the same failures accumulate in the dark.

The engineering world arrived at the same conclusion and built a practice around it: the **blameless postmortem**. After a failure, teams at companies like Google conduct a written review whose explicit ground rule is that no individual is named as the cause. The question is never “who made the mistake” but “what about our systems, tools, and conditions made this mistake possible — and likely” (Google SRE). The premise, articulated early by the engineer John Allspaw and his colleagues at Etsy, is that the people involved were competent professionals acting reasonably given what they knew at the time, and that if a single human error could bring the system down, the deeper fault lies with a system that allowed it to (Allspaw, 2012). Re-

move the fear, and people tell you exactly how the failure happened — which is the only way to stop it recurring.

ACCOUNTABILITY IS NOT BLAME

Blamelessness is the most misunderstood idea in this chapter. It does not mean no one is responsible and anything goes. The safety scientist Sidney Dekker draws the distinction sharply through the idea of a *just culture*: an atmosphere of trust in which people are encouraged and even rewarded for providing essential safety information, but in which there remains a clear line between acceptable behaviour and the genuinely reckless or malicious kind. Blame is backward-looking and asks who should pay. Accountability is forward-looking and asks who will act to change what went wrong. A just culture holds people accountable for their conduct while refusing to let the search for a culprit substitute for the harder work of fixing the system that shaped it.

Dekker frames the whole shift as a move from an *old view* of human error to a *new* one. In the old view — the “bad apple” theory — human error is the cause of trouble: systems are basically safe, and unreliable people make them dangerous, so the fix is to find and remove the bad apples. In the new view, human error is not a cause but a *symptom* — a consequence of deeper trouble designed into the system, the place where well-meaning people collide with tools, pressures, and conditions stacked against them. The two views lead to opposite responses. The old one hunts for someone to discipline and learns nothing. The new one treats every error as a window into the system and uses it to make the system better. Which brings us to the organ-

izations that have made this their entire way of operating — and almost never fail as a result.

CHAPTER TEN

The Mindful Organization

Some organizations operate in conditions so unforgiving that a single mistake means catastrophe — and yet they almost never have one. They achieve it not by eliminating error but by becoming relentlessly, collectively obsessed with it.

An aircraft carrier flight deck is one of the most dangerous workplaces ever devised: jets landing and launching feet apart, in darkness, on a pitching steel platform, amid jet blast and arresting cables under lethal tension. By any reasonable expectation it should be a place of frequent disaster. It is not. Carrier decks, along with nuclear power plants and air-traffic-control centres, belong to a class the organizational theorists Karl Weick and Kathleen Sutcliffe call **high-reliability organizations** — outfits that face overwhelming potential for catastrophe yet manage to operate for years almost without one. Studying how they do it, Weick and Sutcliffe identified a shared posture they call *collective mindfulness*, expressed in five principles (Weick & Sutcliffe, 2007).

The first is a **preoccupation with failure**. Far from relaxing when things go well, these organizations treat every small lapse, anomaly, and near-miss as a symptom of something that

could become catastrophic, and they investigate it as such. A near-miss is not “we got away with it” but a free lesson in how the system can fail — data that, in a complacent organization, would be discarded as a non-event. The second is a **reluctance to simplify**. They resist the tidy, reassuring explanation and deliberately preserve the messy nuances and minority interpretations that a simpler story would discard, because the discarded detail is often where the next failure hides.

The third is **sensitivity to operations** — sustained attention to the actual, granular state of work at the front line, rather than to an idealized picture of it drawn in a distant office. The fourth is a **commitment to resilience**: building the capacity to detect, contain, and recover from the errors that will inevitably slip through, rather than trusting prevention alone. And the fifth is **deference to expertise**: when a crisis hits, decision-making authority migrates to whoever has the relevant knowledge, regardless of their rank, instead of freezing at the top of the hierarchy. Notice that none of the five is about preventing all error. Every one is about *seeing* error early, taking it seriously, and staying able to respond.

How the holes line up

To understand why mindful organizations focus on systems rather than individuals, it helps to picture how accidents actually happen. The psychologist James Reason gave us the enduring image: the **Swiss cheese model**. An organization’s defences against failure are like a stack of slices of Swiss cheese. Each slice — a procedure, an alarm, a checklist, a second pair of eyes — has holes in it, the inevitable weaknesses of any single

safeguard. On a normal day the holes do not align, and a hazard that slips through one layer is caught by the next. Disaster strikes only on the rare occasion when the holes in every layer momentarily line up, opening a clear path from hazard to harm (Reason, 1990).




Figure 10.1. Reason's Swiss cheese model. No single safeguard is whole; each has gaps. Harm occurs only when gaps across every layer align. The lesson is systemic: chasing the last person in the chain ignores the holes — and the holes are the real cause.

Reason's model carries a second distinction that reframes blame entirely. The holes come in two kinds. *Active failures* are the unsafe acts at the sharp end — the slip, the misread dial, the wrong button — committed by the person closest to the accident, who is also the easiest to blame. *Latent conditions* are the holes built into the system long beforehand by decisions made far from the front line: understaffing, bad interface design, production pressure, missing safeguards. Latent conditions can lie dormant for years before they combine with an active failure to cause harm. Blaming the operator at the sharp end is satisfying and almost

always misses the point, because the latent conditions — the ones that were going to catch *someone* eventually — remain untouched, waiting for the next person.

When a field learns to count its dead

No domain has confronted this more publicly than medicine. In 1999, the U.S. Institute of Medicine published *To Err Is Human*, which estimated that between 44,000 and 98,000 Americans died each year from preventable medical errors in hospitals — and, crucially, argued that the cause was not an epidemic of bad doctors but badly designed systems that set competent people up to fail (Kohn, Corrigan, & Donaldson, 2000). The report's power was partly in its title: to err is human, so stop trying to perfect the humans and start engineering systems that make errors less likely and less lethal. A later analysis by Martin Makary and Michael Daniel put the toll far higher, at around 250,000 deaths a year — enough, they argued, to make medical error the third-leading cause of death in the country (Makary & Daniel, 2016).

And here the book's own discipline demands a caveat, because both numbers are contested — the higher one sharply so, criticized as an over-extrapolation from small samples, including by researchers in patient safety itself. This is not a weakness in the argument; it is the argument. Even our estimates of how often we err are themselves error-prone and subject to revision. What is not in dispute is the reframing. The journalist Matthew Syed, in *Black Box Thinking*, contrasts two professional cultures: aviation, which after every accident pulls the black box, investigates without blame, and publishes the lessons so that every airline on earth corrects the same flaw — and a medical culture

that long treated error as private shame to be buried. The difference between them is not the quality of the people. It is whether the system treats failure as evidence to be studied or as guilt to be hidden. That choice — and the protocols that follow from it — is what the final part of this book is about.

PART FOUR

IV

The Practice of Correction

From understanding to method. The concrete protocols that turn error into improvement, the culture that makes people use them honestly, and the oldest and most powerful error-correcting machine humanity has ever built.

CHAPTER ELEVEN

The Error-Correction Protocol

Good intentions toward failure dissolve under pressure. What survives is structure — a small set of repeatable routines that force errors into the open and convert them into a decision about what to do differently.

Everything established so far — that error is information, that surprise teaches, that safety surfaces mistakes, that systems beat blame — remains inert until it is built into a routine. Under deadline and ego, people do not rise to their good intentions; they fall to their habits. The practical question of this book is therefore narrow and answerable: what habits reliably turn a mistake into a lesson? A handful of well-tested protocols do exactly that, and each one is a way of forcing the error-correction loop to run when it otherwise would not.

The simplest is the **after-action review**, developed by the U.S. Army as a disciplined debrief after any significant action. It asks four plain questions: What was supposed to happen? What actually happened? Why was there a difference? And what will we do differently next time? The power is in the structure. By separating the intended outcome from the actual one and dwelling on the gap — the prediction error, made institutional — it converts a vague sense that things went sideways into a specific, documented, actionable lesson. Done routinely and without

blame, it turns every operation, successful or not, into a source of learning.

The autopsy before the death

The after-action review extracts lessons after the fact. Its mirror image extracts them before any damage is done. The psychologist Gary Klein devised the **premortem**, a technique he laid out in a 2007 article that has since spread through boardrooms and project teams (Klein, 2007). At the start of a project, when the plan is set and optimism is high, the team is told to imagine that it is months in the future and the project has already failed — catastrophically, unmistakably. Each person then writes down, independently, every reason they can think of for why it failed. Only afterward do they share, round-robin, and consolidate.

The genius of the premortem is psychological. In a normal planning meeting, raising doubts feels like disloyalty, so the people who quietly foresee problems stay silent. By stipulating that the failure has *already happened*, the premortem reframes pessimism as insight and gives the quiet skeptic permission to speak: not “I worry this might fail,” which sounds like a lack of faith, but “here is why it *did* fail,” which sounds like analysis. It rests on a documented effect called *prospective hindsight*: research by Mitchell, Russo, and Pennington found that imagining an event has already occurred, rather than that it might, increases people’s ability to identify reasons for the outcome by roughly thirty percent (Mitchell, Russo, & Pennington, 1989). The exercise takes twenty or thirty minutes and is among the highest-value-per-minute moves a team can make. Its one com-

mon ruin is worth flagging: run it with the most senior executive in the room speaking first, and the candour it depends on collapses.

KEEP A DECISION JOURNAL

The individual analogue of these team rituals is a decision journal, and it exists to defeat one specific enemy: hindsight bias. Before a consequential decision, write down what you expect to happen, why, and how confident you are. Then, crucially, do nothing with it until the outcome arrives. Reviewing the journal later does two things no memory can. It separates the quality of your *decision* from the quality of the *outcome* — a good decision can yield a bad result and vice versa, and only the written record lets you tell them apart. And it confronts you with what you actually believed beforehand, before hindsight rewrote it into something you “knew all along.” A decision journal is a tool for catching yourself in the act of not having known.

One loop, every scale

Step back from the particular protocols and the same shape appears in all of them. The dopamine neuron of Chapter Three predicts, acts, compares its prediction to reality, and updates. The deliberate practice of Chapter Six attempts, fails, gets feedback, and adjusts. The after-action review states the intent, observes the result, examines the gap, and changes the plan. These are not analogies. They are the identical loop — predict, act, compare, update — running at the scale of a synapse, a skill, and

an institution. To build a learning protocol is simply to make that loop explicit, deliberate, and impossible to skip.




Figure 11.1. The error-correction loop. Make a prediction, act, compare the result to what you expected, update — then predict again. A neuron runs it in milliseconds; a team runs it across a quarter. Every protocol in this chapter is a way of refusing to let the loop break.

Protocols, though, are only as good as the willingness to run them honestly, and that willingness cannot be mandated into existence. A team handed a premortem template but punished for the doubts it surfaces will learn to fill in the form and lie. The protocols are the skeleton; what animates them is a culture in which surfacing error is genuinely safe and genuinely valued — which is the harder thing to build, and the subject of the next chapter.

CHAPTER TWELVE

Building a Learning Culture

Tools do not create honesty; culture does. And culture is built less by what is written in the handbook than by what happens, in public, the moment someone admits they were wrong.

A learning culture is not a poster, a value statement, or a mandatory training. It is the accumulated memory of what actually happened to the last several people who brought bad news. If they were thanked, listened to, and helped to fix the problem, the culture learns that error is safe to surface. If they were blamed, sidelined, or made an example of, the culture learns the opposite — and no protocol, however well designed, will overcome that lesson. Culture is taught by consequence, and the consequences that teach it most loudly are the ones attached to admitting a mistake.

This places the heaviest responsibility on whoever holds power, because their reaction to failure sets the price of honesty for everyone below them. The single most effective thing a leader can do to build a learning culture is also the most uncomfortable: to model fallibility openly — to say, in front of the team, “I got this wrong, here is what I missed, here is what I am changing.” When the most senior person in the room treats their own error as ordinary and instructive rather than shameful,

they license everyone else to do the same. When they defend their record at all costs, they teach everyone that the way to survive is to do likewise, and the organization goes quiet exactly where it most needs to speak.

What you measure, and what you reward

Metrics shape behaviour, and naive metrics around error are actively dangerous. An organization that measures teams by the number of reported errors — and punishes the high numbers — has not created safety; it has created an incentive to stop reporting. As Amy Edmondson's hospital research showed, a low reported-error rate can mean a frightened team rather than a careful one. The healthier instinct is to reward the *reporting* and the *response*: to treat a surfaced near-miss as a contribution, to make reporting effortless and visibly consequential, and to measure whether problems, once raised, actually get fixed. Aviation institutionalized exactly this with confidential, non-punitive reporting systems that let crews flag hazards without fear — turning a flood of small admissions into a continuously improving safety record.

A low error count can mean a careful team or a frightened one. The number alone cannot tell you which, and treating it as virtue guarantees the second.

The same logic argues for making the blameless review routine rather than reserving it for catastrophes. If the only time a team

dissects a failure is after a disaster, the exercise carries an unavoidable charge of crisis and blame. If reviewing what went wrong is simply what the team does after everything — a normal, frequent, low-stakes habit — then it loses its sting and becomes what it should be: ordinary maintenance of how the work is done. Routine drains the fear; fear is what stops the learning.

BEWARE THE SLOGAN

The language a culture uses about failure matters, and the fashionable slogans corrode as often as they help. “Fail fast, fail often” was meant to license intelligent experiment, but stripped of its conditions it excuses sloppiness and celebrates failure for its own sake — including the basic, preventable failures that diligence should have caught. The honest version is harder to fit on a wall: fail *well*. Prevent the failures that careful work prevents, design experiments whose failures are informative, and extract the lesson from every one. The goal was never to fail more. It was to waste fewer of the failures you were always going to have.

A learning culture, then, is the human infrastructure that makes the protocols of the previous chapter run on truth instead of theatre: leaders who go first in admitting error, metrics that reward candour over the appearance of perfection, and the steady normalization of looking honestly at what went wrong. Built well, it turns an organization into something that compounds — a system that gets steadily better precisely because it refuses to hide from its mistakes. And that capacity, scaled up beyond any single institution, turns out to be the engine behind

the greatest knowledge-generating enterprise our species has ever assembled.

CHAPTER THIRTEEN

Science, the Engine of Correction

Science is not a body of facts but a method for being wrong productively — humanity’s institutionalized machine for finding and removing its own errors. It is the error-correction loop, scaled to a civilization.

We tend to think of science as a collection of established truths. It is better understood as a procedure for systematically eliminating falsehoods. The philosopher Karl Popper located the dividing line between science and everything else not in proof but in *refutability*: a claim is scientific only if it forbids something, only if there is some observation that could prove it wrong. A theory that explains every possible outcome — that can never be contradicted by any evidence — tells us nothing, precisely because it takes no risk. Real knowledge sticks its neck out. It makes predictions that reality could falsify, and earns its standing by surviving the attempts (Popper, 1963).

This inverts the popular picture of how science works. We do not, on Popper’s account, prove theories true and bank them. We propose bold conjectures and then attack them, and we keep the ones that have so far resisted every refutation — always provisionally, always open to being overturned by the next test. Knowledge advances not by accumulating confirmations but by

the *elimination of error*: each refuted theory clears the ground for a better one. Science is, in its deepest logic, an error-correction process wearing the robes of discovery. Being proven wrong is not a breakdown of the enterprise. It is the enterprise working.

The machinery of organized doubt

What makes science uniquely powerful is that it does not trust any individual to overcome their own biases — the hindsight, the confirmation-seeking, the self-serving distortions of Chapter Eight — and instead builds the correction into the social structure. Peer review invites others to find the flaws an author cannot see. Replication tests whether a result was real or a fluke. Controlled experiment isolates cause from coincidence. Open data and methods let strangers check the work. Each of these is an external error-correcting mechanism, compensating for the fact that no mind can reliably audit itself. Science is organized distrust, turned to constructive ends.

The machinery is imperfect, and its imperfections prove the point rather than undermining it. In recent years psychology and other fields have confronted a *replication crisis* — the discovery that many published findings do not hold up when the experiments are run again. It would be easy to read this as a failure of science. It is the opposite: it is science applying its own error-correction to itself, catching its own mistakes, and tightening its methods in response. The debates we met earlier in this book — over how much deliberate practice really explains, over how large the effects of mindset truly are — are not embarrassments to be hidden. They are the system doing precisely

what it is built to do, revising its estimates in public as better evidence arrives. A field that can say “we were overconfident, and here is the corrected figure” is not weak. It is the only kind of field worth believing.

*A field that can announce its own errors
and correct them in public is not failing. It
is the only kind worth believing.*

It is no coincidence that the same loop appears at the frontier of our newest technology. The artificial-intelligence systems now reshaping the world learn, at bottom, the way Schultz’s neurons did — by making predictions, measuring how wrong they were, and adjusting to reduce the error next time. The reinforcement-learning algorithms driving these systems are direct descendants of the temporal-difference rule the brain was already running. From the synapse to the scientific community to the machine, the pattern holds without exception. Learning, wherever it occurs and at whatever scale, is the detection and correction of error. There is no other kind.

What the loop asks of us

Which returns us, finally, to where the book began — to the flinch in the word *failure*, and the verdict it seems to pass. The science assembled across these pages dissolves that verdict. To be wrong is not a deficiency to be ashamed of; it is the universal precondition of finding out anything at all, the signal that carries every bit of information worth having, the raw material

from which skill, knowledge, and reliability are built. The neuron, the expert, the safe team, the mindful institution, and the scientific community do not differ in whether they err. They all err. They differ in one thing only: how well they have learned to notice it, face it, and correct it.

That is a discipline, and like any discipline it can be practised and improved — in a life, a team, or a civilization. It asks for the humility to write down a prediction that might be wrong, the courage to look at the result when it is, the honesty to name the mistake without flinching, and the structure to make sure the lesson is not lost. The failure was never the thing to fear. The thing to fear is the error you refuse to discover — the one you defend, or hide, or explain away, and so are condemned to repeat. To live and work well is not to stop being wrong. It is to become, deliberately and without shame, very good at being wrong: to treat every error as what it has been all along, the most useful thing that can happen to you.



CODA

The Discipline of Being Wrong Well

A short creed to carry out of the book and into the work.

Everything in these pages reduces to a single change of posture toward a single moment: the instant you discover you were wrong. The flinch tells you to look away, defend, explain, or hide. The science tells you to lean in, because that instant is not the failure — it is the arrival of the only information that can make you better. The whole discipline is built on choosing the second response when every instinct urges the first.

So make your predictions where they can be checked. Write the number down. Seek the difficulty rather than the ease, because ease is the feeling of learning nothing. When the result comes back against you, read it as a measurement, not a verdict. Build the loop into your habits — the review, the journal, the premortem — so that finding your errors does not depend on a mood you may not be in. And if others answer to you, make it safe for them to be wrong out loud, because the alternative is not fewer mistakes but mistakes you never get to see.

None of this makes error rare. Nothing does, and nothing should; a life or an organization with no errors is one that has stopped reaching past what it already knows. The aim was never a record without failures. It was to stop wasting them — to extract, from each thing that goes wrong, the correction it was

carrying all along. That is what the neuron does in milliseconds, what the expert does across years, what the safe team does after every operation, and what science does across centuries. It is available to you, too, starting with the next mistake you are honest enough to notice.

Be wrong. Then be a little less wrong. Repeat, without shame, for the rest of your life. There is no other path to anything worth knowing.



REFERENCE

Glossary of Key Terms

The core vocabulary of error and its correction, as used throughout this book.

Prediction error

The gap between what was expected and what actually occurred. The fundamental quantity from which learning is computed, at every scale from a single neuron to a scientific community.

Reward prediction error

The brain's teaching signal, carried by dopamine: a burst when an outcome is better than predicted, no change when it matches, a dip when it is worse. It encodes surprise about reward, not reward itself.

Temporal-difference learning

The machine-learning rule, developed by Sutton and Barto, in which an agent improves by adjusting its predictions in proportion to its prediction errors. The same rule the dopamine system turned out to implement.

Error-related negativity (ERN)

A sharp electrical signal generated by the anterior cingulate cortex within about fifty milliseconds of a mistake — the brain flagging its own error, often before conscious awareness.

Predictive processing

The framework in which the brain is fundamentally a prediction machine: perception works by generating expectations and propagating upward only the error, the mismatch between prediction and sensory input.

Desirable difficulty

An obstacle introduced into learning that slows acquisition and worsens immediate performance, yet improves long-term retention and transfer — for example, spacing, interleaving, or retrieval practice.

Retrieval practice

Strengthening memory by struggling to recall information rather than rereading it. A leading example of a desirable difficulty; effective even when the retrieval attempt initially fails, provided feedback follows.

Deliberate practice

Effortful, goal-directed practice at the edge of current ability, guided by immediate feedback, with errors identified and corrected — as distinct from mere repetition of what one can already do.

Growth mindset

The belief that abilities can develop through effort and strategy, which shapes whether a person treats a mistake as information or as a verdict. Its real effects are modest and concentrated among at-risk learners.

Hindsight bias

The “knew-it-all-along” effect: once an outcome is known, we misremember having expected it, erasing the genuine prior uncertainty and, with it, the lesson.

Survivorship bias

The error of drawing conclusions only from the cases that “survived” a selection process while overlooking the failures that are no longer visible — named for Abraham Wald’s analysis of returning bombers.

Psychological safety

A shared belief that a team is safe for interpersonal risk-taking — that one can admit a mistake, ask a question, or raise a concern without humiliation or punishment. The precondition for surfacing error.

Blameless postmortem

A structured review of a failure whose ground rule is to seek the systemic and contributing causes rather than to identify an individual culprit, on the premise that fear produces concealment.

Just culture

An organizational climate of trust that encourages reporting of safety-critical information while maintaining a clear line between acceptable behaviour and the genuinely reckless. Accountability without reflexive blame.

Latent condition

In Reason’s Swiss cheese model, a weakness built into a system long before any accident — poor design, understaffing, production pressure — that lies dormant until it aligns with an active failure to cause harm.

High-reliability organization

An organization that operates in hazardous conditions yet rarely fails, through “collective mindfulness”: preoccupation with failure, reluctance to simplify, sensitivity to operations, commitment to resilience, and deference to expertise.

Premortem

A planning exercise in which a team imagines a project has already failed and works backward to identify why, exploiting *prospective hindsight* to surface risks that ordinary caution leaves unspoken.

Falsifiability

Popper's criterion for a scientific claim: that it forbids something and could, in principle, be proven wrong by observation. Knowledge advances by the elimination of error, not the accumulation of proof.

SOURCES

References

Every empirical claim and figure in this book is attributed below and was checked against primary or authoritative sources at the time of writing.

ON ERROR AND ITS KINDS

- Box, G. E. P. (1976).** Science and statistics. *Journal of the American Statistical Association*, 71(356), 791–799.
- Edmondson, A. C. (2011).** Strategies for learning from failure. *Harvard Business Review*, 89(4).
- Edmondson, A. C. (2023).** *Right Kind of Wrong: The Science of Failing Well*. Atria Books.
- Korzybski, A. (1933).** *Science and Sanity: An Introduction to Non-Aristotelian Systems and General Semantics*. Institute of General Semantics.
- Reason, J. (1990).** *Human Error*. Cambridge University Press.

THE BRAIN THAT LEARNS FROM ERROR

- Clark, A. (2013).** Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36(3), 181–204.
- Frank, M. J., Woroch, B. S., & Curran, T. (2005).** Error-related negativity predicts reinforcement learning and conflict biases. *Neuron*, 47(4), 495–501.
- Gehring, W. J., Goss, B., Coles, M. G. H., Meyer, D. E., & Donchin, E. (1993).** A neural system for error detection and compensation. *Psychological Science*, 4(6), 385–390.
- Holroyd, C. B., & Coles, M. G. H. (2002).** The neural basis of human error processing: Reinforcement learning, dopamine, and the error-related negativity. *Psychological Review*, 109(4), 679–709.

- Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, 275(5306), 1593–1599.
- Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27, 379–423.
- Sutton, R. S., & Barto, A. G. (2018). *Reinforcement Learning: An Introduction* (2nd ed.). MIT Press.

PRACTICE, DIFFICULTY, AND MINDSET

- Bjork, E. L., & Bjork, R. A. (2011). Making things hard on yourself, but in a good way: Creating desirable difficulties to enhance learning. In *Psychology and the Real World* (pp. 56–64). Worth.
- Dweck, C. S. (2006). *Mindset: The New Psychology of Success*. Random House.
- Ericsson, K. A., Krampe, R. T., & Tesch-Römer, C. (1993). The role of deliberate practice in the acquisition of expert performance. *Psychological Review*, 100(3), 363–406.
- Gladwell, M. (2008). *Outliers: The Story of Success*. Little, Brown.
- Kornell, N., Hays, M. J., & Bjork, R. A. (2009). Unsuccessful retrieval attempts enhance subsequent learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 35(4), 989–998.
- Macnamara, B. N., Hambrick, D. Z., & Oswald, F. L. (2014). Deliberate practice and performance in music, games, sports, education, and professions: A meta-analysis. *Psychological Science*, 25(8), 1608–1618.
- Mueller, C. M., & Dweck, C. S. (1998). Praise for intelligence can undermine children's motivation and performance. *Journal of Personality and Social Psychology*, 75(1), 33–52.
- Sisk, V. F., Burgoyne, A. P., Sun, J., Butler, J. L., & Macnamara, B. N. (2018). To what extent and under which circumstances are growth mind-sets important to academic achievement? Two meta-analyses. *Psychological Science*, 29(4), 549–571.
- Yeager, D. S., et al. (2019). A national experiment reveals where a growth mindset improves achievement. *Nature*, 573, 364–369.

BIAS, SAFETY, AND SYSTEMS

- Allspaw, J. (2012). Blameless PostMortems and a Just Culture. *Code as Craft* (Etsy Engineering Blog).
- Dekker, S. (2014). *The Field Guide to Understanding 'Human Error'* (3rd ed.). Ashgate.
- Edmondson, A. C. (1999). Psychological safety and learning behavior in work teams. *Administrative Science Quarterly*, 44(2), 350–383.
- Fischhoff, B. (1975). Hindsight $\neq$ foresight: The effect of outcome knowledge on judgment under uncertainty. *Journal of Experimental Psychology: Human Perception and Performance*, 1(3), 288–299.
- Google (re:Work). Guide: Understand team effectiveness (Project Aristotle). rework.withgoogle.com.
- Google SRE. Postmortem culture: Learning from failure. In *Site Reliability Engineering*. sre.google.
- Kohn, L. T., Corrigan, J. M., & Donaldson, M. S. (Eds.). (2000). *To Err Is Human: Building a Safer Health System*. Institute of Medicine. National Academies Press.
- Makary, M. A., & Daniel, M. (2016). Medical error—the third leading cause of death in the US. *BMJ*, 353, i2139.
- Mangel, M., & Samaniego, F. J. (1984). Abraham Wald's work on aircraft survivability. *Journal of the American Statistical Association*, 79(386), 259–267.
- Syed, M. (2015). *Black Box Thinking: Why Some People Never Learn from Their Mistakes—But Some Do*. Portfolio.
- Weick, K. E., & Sutcliffe, K. M. (2007). *Managing the Unexpected: Resilient Performance in an Age of Uncertainty* (2nd ed.). Jossey-Bass.

PROTOCOLS AND THE LOGIC OF SCIENCE

- Klein, G. (2007). Performing a project premortem. *Harvard Business Review*, 85(9), 18–19.
- Mitchell, D. J., Russo, J. E., & Pennington, N. (1989). Back to the future: Temporal perspective in the explanation of events. *Journal of Behavioral Decision Making*, 2(1), 25–38.

Popper, K. R. (1963). *Conjectures and Refutations: The Growth of Scientific Knowledge*. Routledge.

COLOPHON

About the Author

Marie-Soleil Seshat Landry is the founder of Landry Industries, an independent research and publishing venture based in Moncton, New Brunswick, Canada. Her work spans ethical open-source intelligence, bio-based materials research, and the governance of artificial intelligence, united by a single preoccupation: how knowledge can be produced honestly, and how systems — human and machine alike — can be built to find and correct their own errors.

That preoccupation shaped this book and its method. Every empirical claim was retrieved, read, and attributed; contested findings were reported as contested rather than smoothed into false certainty; and the line between what is established and what is provisional was kept visible throughout. It is the same discipline the book argues for, applied to the book itself: **retrieve, store, cite** — and remain willing to be corrected.

She publishes across an extensive open archive under **ORCID 0009-0008-5027-3337**.