THE INTERIOR ARCHITECTURE OF TRANSFORMATION

# The Logic Trap

*How your mind builds the case against itself*

JOSHUA T. FRASER, ED.D.

THE INTERIOR ARCHITECTURE OF TRANSFORMATION

# The Logic Trap

CONTENTS

INTRODUCTION

# The Perfect Process

The conference room had that particular stillness that comes when everyone knows a decision matters. Eight of us around the table: three principals, the HR director, two instructional coaches, a teacher union representative, and me, chairing the committee. We were hiring an equity coordinator. Not a soft position. This person would challenge practices, facilitate difficult conversations, and hold mirrors up to people who didn't want to see what was reflected. The wrong hire wouldn't just fail; they'd set the work back years.

Two finalists remained. I'll call them Melissa and Terri.

Melissa was polished. She arrived early, made eye contact with everyone individually, referenced our strategic plan by page number. When she talked about equity, she used the language: systemic barriers, cultural responsiveness, intersectionality. She cited Ladson-Billings. She mentioned Kendi. She spoke about her previous district's journey with exactly the right mix of pride and humility.

Terri was different. She paused before answering questions, long pauses that made the room shift slightly. Her language was less fluent, more searching. She didn't cite theorists; she told stories. When asked about systemic barriers, she described sitting with a mother who couldn't get her son's IEP meeting scheduled during hours she could attend.

By the time we finished the interviews, I knew who I wanted to hire.

I ran the debrief carefully. Everyone shared their observations. I asked follow-up questions. I invited dissenting views. I summarized points of consensus. Here's what I didn't notice: when someone praised Melissa, I nodded and asked them to say more. When someone raised a concern about her, I asked "clarifying questions" that reframed the concern as a positive. When the union representative said Terri seemed more authentic, I wondered aloud whether she had "the polish for cabinet-level work."

I was running a process I thought was fair. I was actually running a machine that sorted for mirrors.

We hired Melissa. The vote was 6-2, which felt like strong consensus. Eight months later, Melissa resigned. Not because she was incompetent; she was quite capable. But teachers found her condescending. Her facilitation of equity conversations felt like performance. She talked about equity the way I talked about equity. And that, it turned out, wasn't what we needed.

We needed Terri. And I had built a process that made hiring her impossible.

This book exists because I needed it. Not in the abstract. I needed it that day in the conference room, running a process I thought was fair while my thinking quietly poisoned every decision I made. Something in me recognized Terri as the stronger candidate; the pauses that unsettled the room were doing exactly what an equity coordinator would need to do, sitting with discomfort instead of performing fluency. I felt that recognition in my body during her interview, a sense that her searching, unpolished answers carried a weight Melissa's prepared responses did not. But the recognition never made it into my analysis. By the time I ran the debrief, I was operating from the analytical framework alone, sorting candidates by criteria my thinking had constructed to confirm what it had already decided. The felt recognition and the analytical process were running on separate tracks, and the analytical track was louder, more articulate, more rewarded by the institutional culture of the room. I'd read about cognitive biases. I could name them in others. I'd even taught workshops on recognizing distorted thinking. None of that helped, because knowing about bias didn't protect me from it. The name of the trap doesn't spring you from it. What I've learned since: catching your own distortion is a specific practice of noticing when your thinking has separated from something you are sensing but have not yet examined, and interrupting the separation before it calcifies into a decision.

The framework is deliberately simple. Thinking errors, at least the ones this book addresses, cluster around three failure points I call the grooves: the evidence groove, where we see what we expect to see and confirmation bias curates our information before analysis begins; the attribution groove, where we explain what we see by reaching for the wrong cause and deficit ideology locates problems in people rather than systems; and the prediction groove, where we mistake our assumptions for analysis and construct futures that protect our existing models from scrutiny. These three interact and reinforce each other, masquerading as rigor, as responsibility, as righteousness so effectively that a leader can run all three

CHAPTER 1

# The Brilliant Blunder

Consider Dr. Elena Vasquez. Elena is a superintendent in a mid-sized suburban district with a doctorate in educational leadership, fifteen years of administrative experience, and a genuine commitment to equity. She reads research, attends conferences, and can cite Freire and Ladson-Billings in conversation. By every measure she is a thoughtful, informed leader.

Her district has a persistent problem: significant achievement gaps between white students and students of color, particularly in mathematics. The gaps have held steady for years, and Elena has made closing them a strategic priority. She commissions a comprehensive data analysis. Her research team disaggregates everything: test scores, course enrollment, discipline referrals, attendance patterns. They produce a detailed report identifying where the gaps are widest, which subgroups are most affected, which schools show the largest disparities.

Elena studies the report carefully. She notices a pattern: the gaps are largest in schools with the highest concentrations of low-income students of color. She notices another pattern: those same schools have the highest rates of teacher turnover. A third pattern: teachers in those schools report lower satisfaction with professional development. The causal story writes itself. Under-resourced schools lead to less experienced teachers, which leads to lower quality instruction, which leads to achievement gaps. The intervention is obvious: invest in teacher retention and professional development at the highest-need schools.

Elena launches a major initiative. She redirects funding, hires instructional coaches, implements mentorship programs for new teachers. She monitors progress quarterly. She communicates the strategy to the board, to families, to staff. Three years later, the gaps haven't budged. The logic was sound. The investment was real. The implementation was monitored. Why didn't it work?

Here's what Elena didn't see. She was right that teacher turnover correlated with achievement gaps, but correlation isn't causation, and what she missed was a third variable: school leadership quality. The schools with the highest turnover also had the weakest principals. Teachers weren't leaving because of inadequate professional development; they were leaving because of poor working conditions created by ineffective leadership. The professional development investment couldn't compensate for leadership that demoralized staff, created chaotic environments, and failed to support new teachers through their crucial first years.

Elena's intervention was logical. It just solved the wrong problem. But the deeper issue is this: Elena's thinking process made it nearly impossible for her to see what she was missing. The data she examined, the patterns she identified, the story she constructed, and the solution she implemented all fit together seamlessly. Each step felt rigorous. The conclusion felt earned. What she didn't realize was that her mind had been filtering reality before she ever sat down to analyze it. Her attention had been selectively drawn to certain patterns, turnover and PD satisfaction, while other patterns, principal effectiveness and working conditions, remained invisible, present in the data but absent from her hypothesis and therefore absent from her attention.

This is the anatomy of the brilliant blunder: a failure of perception that intelligence then elaborates, justifies, and defends. Something in Elena's experience, a felt sense that the gaps were connected to conditions she understood and could address, ran ahead of the analytical process and shaped which evidence her analysis would find. The thinking channel built its case on ground the feeling channel had already chosen, and the seamlessness of the result made the separation between them invisible.

## The Structure of the Blunder

The brilliant blunder has a structure worth examining with some precision because it operates the same way across very different contexts.

First, there's a *plausible pattern*. The leader identifies something real in the data: a correlation, a trend, a gap. This isn't hallucination; the pattern actually exists. Elena really did see higher turnover in higher-poverty schools. The pattern wasn't invented.

Second, there's a *causal story* that explains the pattern. This is where the distortion enters. The leader needs to make sense of what they're seeing, and they construct a narrative: X causes Y causes Z. The narrative feels like analysis, often sounds sophisticated, references research,

invokes frameworks, demonstrates knowledge. But it's a story shaped by the grooves in the leader's thinking, and the feeling that the story is right, the sense of coherence it produces in the body, is part of what prevents the leader from testing it against alternatives.

Third, there's an *intervention that follows logically from the causal story*. If the story is right, the intervention should work. Elena's investment in professional development made perfect sense if teacher quality was the primary driver and professional development could improve teacher quality and improved teacher quality would close gaps. Each "if" is a load-bearing assumption, and any one of them could be wrong.

Fourth, there's *confirming feedback that reinforces the original analysis*. This is the cruelest part. Elena monitored her initiative. She saw teachers attending professional development sessions. She saw retention rates improve slightly. She saw satisfaction scores tick up. It looked like the strategy was working, but she was measuring the intervention rather than the outcome. The feedback she gathered confirmed that she was implementing her plan, not that her plan was addressing the actual problem.

Finally, there's *persistent failure that doesn't update the model*. The gaps don't close, but instead of questioning the causal story, the leader often doubles down: more professional development, better implementation, longer timelines. The model that produced the failure is protected from scrutiny because the leader can always find an implementation flaw to blame. The certainty that the story is right, a certainty that lives as much in the body as in the analysis, makes it nearly impossible to step back and ask whether the story itself is wrong.

This is the structure of the brilliant blunder: systematic error emerging from the way intelligent people process information, devastating in educational leadership where the distance between decisions and their consequences is measured in semesters and school years, long enough for any number of alternative explanations to accumulate.

## Dysrationalia and the Intelligence Trap

Why does this happen to smart people? You'd think intelligence would protect against distorted thinking. If you're good at analysis, shouldn't you be good at catching your own mistakes?

The research says no. In fact, there's evidence that intelligence can make certain errors more likely, because intelligence gives you more tools for constructing convincing explanations. When you're smart, you can build elaborate justifications for conclusions you've reached through flawed processes, cite research that supports your view while not seeking research that challenges it, construct sophisticated arguments that sound rigorous but rest on unexamined assumptions. The smarter the leader, the more elaborate the confirmation architecture, and the more convincing the arguments that keep the groove intact.

Psychologist Keith Stanovich calls this "dysrationalia": the inability to think rationally despite adequate intelligence.^1^ He distinguishes between intelligence, which is raw processing power, and rationality, which is the disposition to use that power well. The distinction is not trivial. Stanovich's research demonstrates that the correlation between IQ and rational thinking is weaker than most people assume, and that high-IQ individuals are not reliably better at overriding the heuristics and biases that produce systematic error. Smart people can be profoundly irrational, and their intelligence often makes their irrationality more dangerous because they can argue themselves into increasingly elaborate versions of wrong.

This pattern shows up in research on what's called "myside bias": the tendency to evaluate evidence and generate arguments that favor your existing beliefs.^2^ Studies show that intelligent people are no better than average at evaluating arguments for positions they oppose. They are, however, substantially better at defending positions they already hold. Intelligence makes you a better lawyer for your current worldview, not a better judge of whether that worldview is accurate.

Elena was brilliant at building the case for her professional development initiative. She could cite the research on teacher quality, point to the data on turnover, construct a plausible chain of causation. What she couldn't do, what her intelligence didn't help her do, was seriously entertain the possibility that she was wrong. The sense of rightness she felt about her analysis, the coherence of the story, the emotional satisfaction of having identified a solvable problem, all of it operated below the level of argument where intelligence works. The feeling came first. The argument came to serve it.

## The Aftermath at Westfield Middle

Two years into his restorative practices initiative, Marcus Thompson had lost eleven teachers. He knew the number because his assistant superintendent had flagged it in a memo, the kind of memo that arrives with polite language and sharp implications. What Marcus didn't know, or couldn't let himself know, was how many of those departures represented exactly the feedback he claimed to want.

Marcus was a principal at a diverse urban middle school who had committed fully to restorative practices as an alternative to traditional discipline. He'd read the literature, attended the training, and genuinely believed that punishment-based approaches harmed students, particularly students of color. He wasn't wrong about this; the research on exclusionary discipline is damning. So he implemented circles, created alternatives to suspension, and hired a restorative coordinator. When teachers sent students to the office, Marcus asked them to consider what they could do differently, to examine their own practices before removing students from class.

Some teachers appreciated this. They felt supported in building relationships with students and experienced fewer disruptions as the culture shifted. Other teachers felt abandoned, felt Marcus was prioritizing students over staff, that their concerns about classroom management were being dismissed, that they were

being blamed for student behavior. A few of these teachers were genuinely resistant to equity work. But some were teachers of color who had seen initiatives come and go, who had been burned by leaders who talked about transformation but didn't provide practical support.

Marcus noticed the resistance and categorized it. In his mind, the teachers who struggled with restorative practices were either committed to punitive approaches that harm students or not yet developed enough to understand the deeper work. He created professional development to help them get there, brought in speakers, shared articles, modeled the practices he wanted to see. Something in him registered that the pushback carried a different quality than simple resistance, that some of the voices objecting were voices he respected, but that registration never made it into his thinking about what to do. The analytical framework he was operating from had one category for people who disagreed with his approach, and the category was "not yet ready." The felt recognition that the situation was more complicated than that ran on a separate track, and the analytical track was louder.

The resistant teachers felt lectured to, and some left while others disengaged. The teachers who remained supportive became his inner circle, and he stopped hearing critiques because he'd surrounded himself with agreement. By the time he received the assistant superintendent's memo, his school had one of the highest turnover rates in the district. He'd lost several Black teachers who felt their concerns were dismissed as "resistance." His restorative program was struggling because he'd lost the staff capacity to sustain it.

What happened to Marcus? His commitment to equity wasn't wrong. His critique of punitive discipline wasn't wrong. His belief in restorative practices wasn't wrong. What was wrong was how he processed disagreement. When teachers raised concerns, Marcus didn't hear feedback; he heard resistance. His mind sorted their objections into categories that allowed him to dismiss them: "That's just discomfort with change," or "That's deficit thinking about students," or "They're not ready for this work yet." Some of those categorizations were accurate, because some teachers really were resistant to equity. But Marcus's thinking process couldn't distinguish between resistance that should be challenged and feedback that should inform his practice, because his commitment to justice had become a filter that blocked valid information. The feeling channel, which registered the complexity of what teachers were telling him, and the thinking channel, which had built a framework that classified all dissent, were operating independently, and the framework was drowning out the signal.

This is what makes the brilliant blunder so insidious for leaders who care about equity. The very frameworks that illuminate oppression can also blind us to our own failures. If you've learned to identify deficit thinking in others, you can use that lens to dismiss anyone who questions your approach. If you've learned to recognize resistance to equity, you can categorize all disagreement as resistance. The tools of liberation can become tools of self-protection, and the mechanism is always the same: a felt certainty that runs ahead of analysis, shapes which evidence the analysis finds, and produces conclusions that feel earned because the

process that generated them was invisible.

Marcus wasn't a bad leader. He was a passionate, committed leader whose thinking had grooves that distorted his perception. And because his distortion aligned with values he held dear, he couldn't see it operating.

## Three Grooves

The grooves in our thinking cluster around three failure points. The first is *evidence*: how we see. Our minds don't passively receive information; they actively filter it. We notice what we expect to notice and miss what we don't, weight evidence that confirms our beliefs more heavily than evidence that challenges them, remember the examples that support our narrative and forget the ones that don't. Elena's evidence groove led her to see teacher turnover and professional development satisfaction while missing principal effectiveness. The data on principal leadership was there, in the complaints teachers made during exit interviews, in the patterns of who left and when, but it wasn't in her hypothesis, so it wasn't in her attention.

The second is *attribution*: how we explain. When we see a pattern, we need a story about why it exists, and our minds gravitate toward certain kinds of explanations, particularly those that locate cause in individual character rather than situational factors. We see student failure and attribute it to student motivation, see teacher struggle and attribute it to teacher skill, see gaps and attribute them to the deficits of those experiencing them. Marcus's attribution groove led him to explain teacher pushback as resistance to equity rather than response to inadequate support. He saw behavior, teachers complaining about his approach, and attributed it to character, that they weren't committed to justice, rather than circumstance, that they didn't have the structures they needed to succeed.

The third is *prediction*: how we forecast. We don't just interpret the present; we construct the future, and our constructions often protect our existing models. We predict that interventions will work, that skeptics will come around, that time will vindicate our approach, while catastrophizing change we don't like and minimizing risks we've decided to take. Both Elena and Marcus fell into prediction grooves. Elena predicted that professional development would improve teaching, which would close gaps, which would prove her analysis correct. Marcus predicted that resistant teachers would either get on board or self-select out, which would strengthen his program, which would prove his approach right. Neither seriously imagined a future where they were fundamentally wrong.

These three grooves interact and reinforce each other: a faulty evidence groove leads to a faulty attribution, which leads to a faulty prediction, which filters future evidence to confirm the original distortion. The system is self-sealing, and at every stage, the bodily registration that the analysis is correct, the bodily experience of coherence, operates alongside the analytical process without informing it. The two tracks run

together in time but separately in function, and the separation is the trap.

## Equity and Distorted Sight

Why does this matter so much for equity work? Because equity requires clear sight, and our thinking naturally distorts sight. Every system of oppression survives by distorting perception: racism teaches white people not to see privilege, classism teaches affluent people not to see advantage, ableism teaches non-disabled people not to see exclusion. The water we swim in is designed to be invisible.

Equity work begins with trying to see what systems have taught us not to see. This is essential, but it's not sufficient, because the same minds that have been trained by oppressive systems are also the minds we use to do equity work. We don't get new minds when we commit to justice. We get the same minds, with the same grooves, now trying to perceive what they were designed to miss.

And here's the trap: the act of committing to equity can create new grooves that are just as distorting. When we learn to see deficit thinking, we become skilled at identifying it in others, and that skill can blind us to our own deficit assumptions. When we learn to recognize resistance to equity, we become skilled at categorizing disagreement, and that skill can prevent us from hearing valid critique. When we learn to analyze systems of oppression, we become skilled at structural explanation, and that skill can excuse us from examining our own failures. The grooves don't care about our values. They operate regardless of what we believe. A leader committed to justice is not immune to confirmation bias. A superintendent fighting for equity is not protected from the fundamental attribution error. A principal implementing restorative practices can still predict the future inaccurately and filter evidence that challenges their narrative.

This is why knowledge of oppression is not enough. You can understand systemic racism perfectly and still make brilliant blunders in your leadership. Understanding is necessary but not sufficient. What's also needed is the skill of catching yourself, of recognizing the shape of your own distortion while there's still time to correct it, and that skill begins with noticing the gap between what you feel about a situation and what your analysis says about it, because the gap is where the groove lives.

## The Limits of Self-Examination

There's a particular challenge here that deserves attention. The tools you'd use to analyze the distortion are themselves distorted. Your reasoning examines your reasoning and finds it sound, because your reasoning is doing the examination. It's like trying to see your own blind spot by looking harder.

This is why knowledge of cognitive bias offers limited protection. You can know all about confirmation bias and still fall into it. You can teach workshops on the fundamental attribution error and commit it in the same afternoon. Knowing the name of the trap doesn't spring you from it. The knowledge stays in the analytical channel while the bias operates in the channel underneath it, the one that shapes which evidence reaches the analysis in the first place.

I want to be honest about something here: I am not sure how far the practices in this book actually reach. I've seen them work. I've also seen smart, well-meaning leaders apply every interruption protocol I know and still miss the thing that was right in front of them. The grooves are deep, and the tools for addressing them are themselves subject to the same cognitive limitations they're designed to overcome. I don't know where the ceiling is, whether the practices I'm about to describe represent 20% improvement in catching distortion or 60%. I know they're better than nothing. I believe they're substantially better. But I can't prove that with the certainty I'd like to feel.

What does help is external interruption. You need structures that challenge your analysis from outside your own head, people who will tell you what you're not seeing, protocols that force you to seek disconfirming evidence, practices that slow down the automatic pattern-matching your brain does and create space for alternative interpretations. You need, in short, to build a system that compensates for the separation between what you sense and what you analyze, one that forces the two channels into the same room.

That's what the rest of this book offers: a set of practices for catching distortion faster, a protocol you can run in real time, in meetings, in data analysis, in difficult conversations, that interrupts the brilliant blunder before it produces its damage. The goal isn't to eliminate the grooves; they're features of human cognition, and they're not going away. The goal is to know them so well that you can feel them operating, and interrupt them before they do their worst.

## What's Ahead

Before we proceed, a word about what's ahead. This work is uncomfortable, and the discomfort is personal rather than intellectual. To catch your own distorted thinking, you have to admit you're capable of it. You have to accept that your careful analysis might be wrong. You have to hold open the possibility that people you've dismissed as resisters might be seeing something you've missed.

For leaders committed to justice, this is particularly hard. We want to be the ones who see clearly. We've built identities around being the people who get it while others don't. Admitting that our thinking has grooves, that our equity analysis might be filtering out valid information, can feel like a betrayal of the work itself.

It's a deepening of the work. The leaders who do the most damage aren't the ones who oppose equity; they're the ones who pursue equity with distorted thinking, who build elaborate initiatives on faulty causal stories, who dismiss legitimate feedback as resistance, who measure their interventions while ignoring their outcomes. They mean well. They work hard. They cause harm.

I've been that leader. More than once. The hiring committee was one example. There are others I haven't told you, and some I probably haven't fully seen yet. The grooves don't disappear when you become aware of them. They just become something you can work with: patterns you recognize, interruptions you can make, corrections you can attempt before the cost is paid by someone else.

That's the work of clear thinking. The recognition that you are always in the process of catching up to your own distortion, that the span between what you sense and what you analyze is the space where the grooves operate, and that closing that distance, even partially, even temporarily, is the practice that makes the difference.

## Confidence and Accuracy

There's a pattern in how the brilliant blunder unfolds that deserves attention: the relationship between confidence and accuracy.

When you know a lot about something, your confidence increases, and this makes sense because expertise should produce confidence. But confidence also increases when you know less than you think you know, and the feeling of certainty doesn't distinguish between actual knowledge and apparent knowledge.

Dr. Jennifer Nakamura studies leadership decision-making. She describes what she calls the "expertise illusion": leaders who have been successful develop confidence in their judgment, and that confidence feels earned because it often is earned. But the same confidence extends to domains where the leader has less expertise than they believe.^3^

"The principal who's been effective at school turnaround develops confidence in their judgment generally," Jennifer explains. "So when they face a curriculum decision, which requires different expertise, they bring the same confidence. But the confidence isn't calibrated to the actual domain. They're confident about something they don't know well."

This is how the brilliant blunder becomes systematic. Your success in one area creates a generalized confidence that extends to other areas, and you stop noticing the boundaries of your expertise because the feeling of knowing is the same everywhere. The certainty that serves you in your area of expertise may betray you outside it, and because the feeling is identical in both cases, the body provides no signal that a boundary has been crossed.

Consider what this means for equity work specifically. Many leaders come to equity with confidence built in other domains. They were effective instructional leaders who ran successful schools and improved outcomes on various metrics. That track record creates confidence, and the confidence transfers. But equity work requires a different kind of expertise: understanding systems of oppression, historical patterns, cultural dynamics, the ability to see what you've been trained not to see. The skills that made you successful as an instructional leader may or may not transfer.

Dr. Thomas Andersen led high-performing schools before taking on equity leadership. "I brought all my confidence with me," he says. "I had a track record. I knew how to improve schools. I figured equity was just another improvement problem."^4^

It took two years for him to recognize his error. "Equity wasn't another improvement problem. It required me to question the very frameworks I had used to define improvement. My confidence was getting in the way. I was so sure I knew how to lead that I couldn't see where I was failing."

The brilliant blunder in equity work often looks like this: a confident leader applies frameworks that have worked before without recognizing that equity requires different frameworks. They see data through lenses developed for other purposes, attribute outcomes to causes that fit their existing mental models, predict that familiar approaches will produce familiar results. The grooves aren't neutral. They were formed in systems that advantage some people over others, and when you bring those grooves to equity work, you risk perpetuating exactly what you're trying to change.

## The Permanence of the Work

I want to return to something I said earlier and push on it.

I said the goal isn't to eliminate the grooves, that they're features of human cognition and they're not going away. But let me be more specific about what this means. The goal isn't acceptance of distortion; the goal is recognizing that the work is never finished. You are building practices that catch distortion faster, and those practices need to be sustained because the grooves persist and regenerate and adapt.

This distinction matters because leaders often approach self-improvement as a problem to be solved: identify the issue, address it, move on. But the grooves don't get solved. They get managed, interrupted, compensated for, and they persist. If you expect to reach a point where you've fixed your thinking, you'll either give up in frustration or convince yourself you've arrived when you haven't, and both outcomes leave the grooves operating unchecked.

The leaders who do this work best are the ones who've accepted its permanence. They expect ongoing distortion. They've built practices that assume they'll keep making mistakes. They measure progress by speed

CHAPTER 2

# The Evidence Groove

Your mind doesn't passively receive information; it constructs it. Before you open a spreadsheet, before you walk into a data meeting, before you read a single test score, your brain has already decided what counts as evidence and what doesn't. The construction is invisible, feels like seeing, operates with the seamlessness of perception itself, and the seamlessness is the problem, because what feels like neutral observation is actually selection shaped by everything you already believe.

Cognitive scientists have a term for this: selective attention. The brain processes roughly eleven million bits of sensory information per second but can consciously attend to approximately forty.^1^ The ratio is staggering. For every piece of information that reaches your awareness, nearly three hundred thousand pieces do not, and something is choosing which forty survive. That something is not random. It is shaped by prior belief, by expectation, by emotional salience, by what psychologists call "top-down processing": the imposition of existing mental models onto incoming data.^2^ The implications for leaders who believe they are "looking at the evidence" are severe. They are looking at a narrow, pre-filtered slice of evidence that their cognitive architecture has already curated to match what they expected to find, and the looking feels neutral because the filtering happens before awareness begins.

This is the evidence groove: the pattern by which your mind selects, weights, interprets, and remembers information in ways that confirm what you already believe. It does not falsify data. It does not require dishonesty or incompetence. It operates below conscious awareness, shaping the information environment before analysis begins. By the time you sit down to "look at the data," the groove has already determined which data you will look at, which questions you will ask, which patterns you will notice, and which you will not. And underneath the analytical process, a physical recognition of what the data should say is running on its own track, guiding the analysis without appearing in it.

The formal study of this phenomenon began in 1960, when psychologist Peter Wason designed an experiment that remains one of the most cited in cognitive science.^3^ Subjects received a three-number sequence, 2, 4, 6, and were told to discover the underlying rule by proposing their own sequences. The experimenter would confirm or deny whether each proposed sequence fit the rule. Most subjects formed a hypothesis immediately: ascending even numbers. They then tested sequences that confirmed it: 8, 10, 12. Yes. 14, 16, 18. Yes. 20, 22, 24. Yes. They announced their rule with confidence. They were wrong.

The actual rule was simply "any ascending numbers." The sequence 1, 3, 5 would have confirmed it, and so would 3, 97, 412. But subjects rarely tested disconfirming sequences. They generated evidence that supported what they already believed and treated the absence of contradiction as proof. Wason called this confirmation bias, and subsequent research has demonstrated its operation across virtually every domain of human judgment.^4^

The pattern operates at multiple levels simultaneously. When people read arguments about controversial topics, they spend more time engaging with arguments that support their existing views and less time with counterarguments.^5^ When evaluating evidence, they apply rigorous scrutiny to findings that challenge their beliefs while accepting confirming evidence with minimal skepticism, a phenomenon Ditto and Lopez term "motivated skepticism."^6^ When recalling past events, they reconstruct memory to align with current beliefs, a process Kahneman describes as the inability to "reconstruct past states of knowledge."^7^ The bias operates in both directions at once: seeking confirmation and avoiding disconfirmation. The result is an information environment that feels comprehensive but is curated by the very beliefs it appears to validate. And the feeling of comprehensiveness, the bodily sensation that you have considered the evidence thoroughly, is itself part of the distortion, because the feeling is generated by the consistency of what the groove selected, not by the completeness of what was available.

Stanovich's work on dysrationalia is central here: the inability to think rationally despite adequate intelligence.^8^ He distinguishes between intelligence, which is raw processing power, and rationality, which is the disposition to deploy that power well. The distinction matters because it overturns a comfortable assumption. Studies show that intelligent people are no better than average at evaluating arguments for positions they oppose, though they are substantially better at defending positions they already hold.^9^ Intelligence makes you a better lawyer for your current worldview, not a better judge of whether that worldview is accurate. The smarter the leader, the more elaborate the confirmation architecture and the more convincing the arguments that keep the groove intact.

Confirmation bias does not operate alone. It is compounded by a family of related distortions, each reinforcing the others in ways that make the evidence groove deeper and harder to escape.

The availability heuristic, identified by Tversky and Kahneman in 1973, causes people to judge the frequency or probability of events based on how easily examples come to mind.^10^ Whatever is vivid, recent, or emotionally charged feels more common than it actually is, which is why people overestimate the likelihood of plane crashes and underestimate the risk of heart disease, and why a single school shooting reshapes national safety policy while chronic absenteeism, affecting millions more students annually, receives a fraction of the legislative attention. In schools, the dramatic incident governs perception while the quiet pattern goes unexamined. A single violent altercation can occupy months of administrative discussion and reshape an entire building's discipline approach while data showing that student mental health crises have increased 300% over five years sits in a report that no one references after the initial presentation, because the crisis lacks the visceral memorability of the violent incident. The available evidence is not the important evidence; it is the memorable evidence, and those are different categories whose conflation is one of the most common errors in educational decision-making.

Anchoring bias compounds the problem from another direction. Tversky and Kahneman demonstrated that the first piece of information encountered in a judgment task disproportionately influences all subsequent reasoning, even when that information is arbitrary.^11^ In one well-known experiment, subjects who spun a wheel landing on the number 65 subsequently estimated the percentage of African nations in the United Nations as significantly higher than subjects whose wheel landed on 10. The number was random, and it still anchored their judgment. In school improvement, whoever frames the problem first sets the anchor: if the initial analysis defines the challenge as "curriculum alignment," subsequent discussion orbits curriculum even when accumulating evidence points elsewhere. A school improvement team that begins by reviewing state test data has anchored itself to test scores, and alternative data sources, student voice, teacher observation, family perspective, get evaluated against the anchor rather than on their own terms. The team can spend an entire year discussing how to move test scores without once questioning whether test scores are the right measure of what they are trying to accomplish.

Hindsight bias, which Fischhoff documented in a landmark 1975 study and Kahneman later expanded, creates the illusion that outcomes were predictable after they occur.^12^ Once you know that an initiative failed, you "remember" having had doubts about it. Once you know a hire didn't work out, you recall the warning signs you "noticed" during the interview. This retrospective reconstruction is not deliberate dishonesty but the brain's automatic reorganization of memory around known outcomes. The consequence for leaders is that they rarely learn from experience in the way they believe they do, because experience is constantly being rewritten to confirm that they knew what they were doing all along. Fischhoff's subsequent research suggests this phenomenon is remarkably resistant to correction, persisting even when subjects are explicitly warned

about it and motivated to avoid it.^13^ The groove doesn't just filter incoming information; it rewrites the past, and in rewriting the past, it prevents the recognition that something felt wrong at the time but was overridden by the analytical process that was already building its case.

In educational leadership, the evidence groove produces three predictable distortions. Each operates at a different stage of the decision process, and together they create an information ecosystem that feels rigorous while systematically excluding the data that matters most.

The first distortion occurs at the point of data collection. What a school chooses to measure reflects what its leaders already believe matters: if the operating theory holds that achievement gaps are driven by teacher quality, the system measures teacher effectiveness ratings; if the theory points to curriculum, it measures curriculum alignment scores; if the theory implicates family engagement, it tracks parent involvement metrics. Each measurement decision makes certain patterns visible and renders others invisible, and the decision about what to measure typically happens long before analysis begins, in the quiet moment when someone designs the dashboard.

Consider what a typical school improvement dashboard displays: attendance rates, discipline referrals, aggregated test scores, graduation rates, course completion percentages. Now consider what it excludes: student sense of belonging, teacher collaboration quality, classroom-level variation in outcomes, student voice data, staff morale, community perception of the school. The dashboard is not inaccurate; it is incomplete, and the incompleteness is not random. The missing data tends to be precisely the data that would complicate the prevailing narrative about what is working and what is not. Research on organizational attention by William Ocasio suggests that what organizations attend to is governed not by importance but by institutional structures that channel attention toward some issues and away from others.^14^ The dashboard is one such structure. It does not reflect reality; it constructs the reality that leaders then "discover" through analysis.

A principal might take pride in a "data-driven" culture, with color-coded indicators and quarterly reviews and carefully formatted slide decks. But when asked how the school measures whether students feel known by adults in the building, there is silence. When asked how they track which teachers students seek out for support, blank stares. When asked what data exists on student experience of belonging, the answer is a single question buried in an annual climate survey with a 40% response rate. The school has data on everything except what decades of research identifies as the strongest predictor of student success: the quality of relationships.^15^ The evidence groove has made academic metrics visible and relational metrics invisible, and because you cannot act on what you do not measure, the school manages academics while relationships atrophy. The felt quality of a school, the thing teachers and students and families can sense in the hallway within minutes of walking in, has no line on the dashboard and therefore does not exist in the analytical frame.

The second distortion operates at interpretation. The same data point supports multiple explanations, and the explanation a leader selects is shaped by prior belief more than by the data itself. When test scores decline, the data is compatible with at least six narratives: instructional quality worsened, the assessment changed, the student population shifted, external stressors increased, curriculum alignment deteriorated, or some combination. Ross and Nisbett's research on the fundamental attribution error demonstrates that people systematically favor dispositional explanations, those pointing to someone's character, ability, or effort, over situational ones, those pointing to context, structure, or circumstance.^16^ A leader who believes teachers are the problem sees instructional failure in declining scores. A leader who believes the system is the problem sees structural barriers. Each interpretation contains partial truth, and the evidence groove determines which partial truth dominates and which gets dismissed as an excuse.

Consider a district that sees a spike in chronic absenteeism among middle school students. The initial interpretation, offered instinctively in a cabinet meeting, is that families are not prioritizing education. The supporting evidence: calls home go unanswered. The proposed solution: a parent education campaign about the importance of attendance. But student-level data, once someone thinks to request it, complicates this narrative considerably. The absent students are not randomly distributed; they cluster in specific periods, with specific teachers. Student interviews do not reveal parental indifference but statements like "I don't feel safe in third period" and "that teacher doesn't like kids like me." Same data point, two entirely different causal stories, and the first was offered without effort while the second required deliberate investigation. That asymmetry is the evidence groove at work: the felt sense that the initial explanation was right, the comfort of a narrative that located the problem outside the institution, operated faster than the analytical process that might have challenged it.

The third distortion concerns whose testimony counts as data. Evidence is not limited to numbers; it includes observation, experience, narrative, and perspective. When leaders decide whose observations to take seriously, they are making evidence decisions with consequences as profound as any statistical analysis. The typical information flow in schools is hierarchical: department chairs report to principals, principals report to superintendents, data teams present to cabinets. At each level, information is filtered, summarized, and smoothed, and the voices that reach senior leadership are systematically different from the voices that do not. Lipsky's research on street-level bureaucracy demonstrates that the people closest to the work, teachers, counselors, front-office staff, possess knowledge that organizational structures routinely prevent from reaching decision-makers.^17^

The students most in need of being heard are precisely the ones absent from advisory councils and focus groups because they have disengaged, have been pushed out, have learned that adults do not listen or that speaking produces consequences rather than change. This creates a feedback loop that research on sampling bias would predict: leaders hear from relatively satisfied constituents, conclude that conditions are acceptable, and remain unaware of how unrepresentative their sample is. The evidence groove makes the positive visible

and the negative invisible through the ordinary mechanics of who shows up, who speaks, and who has learned not to bother.

There is a particular version of the evidence groove that warrants direct examination because it is wrapped in righteous purpose and therefore resistant to critique.

When a leader has learned to see systemic oppression, they have acquired a powerful analytical lens: they can identify how racism operates in seemingly neutral policies, trace how class advantage reproduces itself through funding structures, recognize how ableism is embedded in curriculum design. The lens is real, and the patterns it reveals are real. But a lens is also a filter that highlights certain phenomena and obscures others.^18^ When the equity lens becomes the sole mode of analysis, it produces its own version of confirmation bias, explaining every outcome through the same framework, attributing every failure to the same structural cause, pointing every proposed solution in the same direction.

This becomes visible when equity-focused initiatives fail to produce the intended results. The lens reveals ongoing systemic barriers, which are genuine, but it does not reveal implementation failures, staff confusion, inadequate support structures, or the leader's own gaps in execution. Those remain invisible because the lens was not designed to detect them. The predictable response is to intensify the existing analysis rather than broaden it: more training on systemic oppression, more examination of implicit bias, more critique of deficit thinking, all of which may be necessary and none of which addresses the possibility that the initiative failed for reasons the equity lens cannot detect. Somewhere in the system, a teacher or a team may be producing the outcomes the initiative was designed to achieve, and the questions of what they are doing and what can be learned from studying success rather than exclusively analyzing failure are rendered invisible when the equity lens directs all attention toward systemic explanation.

The danger here is not the lens itself; every analytical framework is partial. The danger is the conflation of the lens with moral commitment, such that questioning the lens feels like questioning the commitment to justice itself. Tetlock's research on "sacred values" demonstrates that when beliefs become entangled with identity and moral conviction, people process challenges to those beliefs not as intellectual disagreements but as moral threats, triggering defensive reasoning rather than open inquiry.^19^ A groove is a groove regardless of its politics. An unexamined groove oriented toward justice will eventually produce the same distortions as any other, selecting confirming evidence and filtering out the data that would complicate the preferred narrative. The politics of the groove do not exempt it from the cognitive science of the groove.

Three disciplines work against the evidence groove. None eliminates it. Each creates friction against the brain's natural tendency toward confirmation.

The first discipline is seeking disconfirming evidence. The question is direct: what would convince you that you are wrong? If you cannot answer with specificity, your belief is not based on evidence; it is protected from evidence. Wason's experiment demonstrated this six decades ago, and the principle has not changed. The discipline requires building time into the analysis process specifically for contrary data: before finalizing any conclusion, the question must be explicit about whether you have searched as hard for evidence against this view as you searched for evidence supporting it. The answer, in practice, is almost always no.

One method operationalizes this effectively. Before finalizing a major decision, write down what specific evidence would change your mind, not "evidence that I'm wrong" but particular data points, particular findings, particular voices that would shift your conclusion. Then assign someone to look for exactly that evidence. The discomfort of inviting contradiction is the point: if disconfirming evidence exists, encountering it before implementation is cheaper than encountering it after; if it does not exist, the search strengthens the original conclusion rather than weakening it. Either outcome is useful. The only useless outcome is never looking.

The second discipline is auditing your data diet. Map your actual information inputs: not the information you believe you receive, but the information that actually reaches you on a regular basis. Who has access to your attention? What voices are present when decisions are made, and whose are structurally excluded? Whose perspective do you encounter only when something goes wrong? Most leaders who conduct this audit honestly discover that their information sources are narrower than they assumed, the same few perspectives and data types and organizational channels all confirming rather than challenging. The narrowness is not chosen deliberately; it accumulates through routine, hierarchy, and the ordinary friction of organizational life. But its effects are the same regardless of intent: a diet that confirms produces conclusions that confirm, and the leader mistakes the consistency of the signal for the validity of the signal.

The third discipline is building adversarial input structures. The brain will not naturally seek disconfirmation, so external systems must force the encounter: assigning a devil's advocate role in planning sessions, requiring a pre-mortem analysis in which the team imagines the project has failed and works backward to identify causes, creating feedback channels that bypass hierarchy so dissenting perspectives reach decision-makers unfiltered. Gary Klein's research on pre-mortem analysis demonstrates that this technique increases the identification of potential problems by 30% compared to standard prospective analysis.^20^ Some leaders institutionalize this further by requiring every major proposal to include a section titled "The Case Against This Plan," authored by someone other than the proposal's advocate. Others conduct anonymous post-decision surveys specifically asking what concerns went unvoiced in the room. These structures do not guarantee that leaders will hear what they need to hear. They increase the probability, and in a domain governed by cognitive bias, increasing probability is the most honest aspiration.

The question that anchors all three disciplines is simple: what am I not seeing?

The question is most urgent when certainty is highest. Confidence often signals not that the analysis is correct but that the evidence groove is operating smoothly, constructing a picture of reality that fits existing models so seamlessly that the construction itself becomes invisible. The question interrupts the process, forces an accounting of the limits of vision rather than an assumption of completeness: what data is absent from this analysis, whose voice has not been heard, what interpretation has not been considered, what would someone who disagrees with this conclusion see in the same evidence?

The evidence groove does not resolve. It does not get permanently interrupted through a workshop, a new dashboard, or a single difficult conversation. It operates continuously because the cognitive architecture that produces it is not a malfunction to be repaired but the architecture itself: the brain's fundamental strategy for managing an environment that contains far more information than consciousness can process.^21^ The eleven million bits per second do not stop arriving. The forty bits of conscious attention do not expand. The separation between what exists and what you perceive remains, and something must fill that gap. What fills it, absent deliberate intervention, is prior belief, running underneath the analytical process on its own track, shaping the evidence the analysis receives before the analysis begins.

The disciplines described here do not fix the architecture. They compensate for it, decision by decision, question by question. The compensation is never complete. The groove re-forms after every interruption. The work is to build the structures, habits, and relationships that surface what your mind would otherwise filter out, and to accept that this work does not end, because the gap between what you sense and what you analyze is structural rather than incidental, a feature of the cognitive architecture rather than a bug that can be patched. The return to the discipline after each lapse, the narrowing of the divide between when the groove fires and when you notice it firing, is the practice. The reduction in that return time is the progress. There is no state of unbiased perception on the other side.

^1^ Wilson, T. D. (2002). *Strangers to Ourselves: Discovering the Adaptive Unconscious*. Harvard University Press.

^2^ Gilbert, D. T. (1991). How mental systems believe. *American Psychologist*, 46(2), 107-119.

^3^ Wason, P. C. (1960). On the failure to eliminate hypotheses in a conceptual task. *Quarterly Journal of Experimental Psychology*, 12(3), 129-140.

^4^ Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology*, 2(2), 175-220.

^5^ Lord, C. G., Ross, L., & Lepper, M. R. (1979). Biased assimilation and attitude polarization. *Journal of Personality and Social Psychology*, 37(11), 2098-2109.

^6^ Ditto, P. H., & Lopez, D. F. (1992). Motivated skepticism: Use of differential decision criteria for preferred and nonpreferred conclusions. *Journal of Personality and Social Psychology*, 63(4), 568-584.

^7^ Kahneman, D. (2011). *Thinking, Fast and Slow*. Farrar, Straus and Giroux.

^8^ Stanovich, K. E. (2009). *What Intelligence Tests Miss: The Psychology of Rational Thought*. Yale University Press.

^9^ Stanovich, K. E., & West, R. F. (2007). Natural myside bias is independent of cognitive ability. *Thinking & Reasoning*, 13(3), 225-247.

^10^ Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, 5(2), 207-232.

^11^ Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157), 1124-1131.

CHAPTER 3

# The Attribution Groove

We have known for half a century that human beings systematically misattribute cause. The research is not ambiguous. Lee Ross named the fundamental attribution error in 1977, documenting our tendency to explain other people's behavior through character while explaining our own through circumstance.^1^ The literature since then has only deepened the finding. When a student fails a class, the adults around that student reach for explanations that locate the failure inside the student or the student's family with a speed and reliability that would be impressive if it weren't so destructive. Motivation, home life, cultural values, parental engagement, self-control: these are the first words out of the room in virtually every data meeting I have ever attended, and I have attended hundreds, and the consistency of the pattern should alarm anyone who has read the research on what actually drives student outcomes.

It should alarm us because the pattern has consequences. When you explain a student's failure as the product of their character or circumstances, you design interventions that target students and families. When you explain it as the product of institutional conditions, you design interventions that change institutions. Only one of these approaches addresses what is actually within the school's power to change, and the attribution groove, operating below the level of conscious choice, reliably steers toward the other. This is not an abstract problem. It is a resource allocation problem, a human dignity problem, and a justice problem, and it has been documented and diagnosed for decades, and it persists because knowing about it doesn't stop it from operating. The knowing stays in the analytical channel. The groove runs underneath.

The attribution groove does not make you see patterns that aren't there. It makes you explain real patterns in ways that protect your existing model of the world. When educators observe a student struggling, their minds automatically construct a causal story, and the story arrives with the speed and certainty of perception itself: the student's family situation, the student's motivation, the student's background. The "why" feels like observation, but it is inference, and the inference is shaped by grooves that were formed in systems designed to locate deficit in particular bodies and particular communities. Richard Valencia (2010) defines deficit thinking as "the idea that students, particularly those from low-income and minority backgrounds, fail in school primarily because of internal deficiencies."^2^ The deficiency might be intellectual, cultural, familial, or characterological, but the location is always the same: inside the student or their circumstances, never inside the institution. The groove selects for these explanations not because educators are callous but because their minds are doing what all minds do, protecting existing models and minimizing cognitive dissonance, and because the explanations that protect the institution feel true in the body with a completeness that preempts further inquiry.

I want to be specific about what I mean by "feels true in the body." When two teachers in a hallway explain a student's failure as a family problem, and a superintendent walking past hears the explanation and feels it land as obvious, that feeling of obviousness is a felt experience with a physical signature: a settling, a release of cognitive tension, a sense that the problem has been adequately explained. The explanation does not require further investigation because it has already satisfied both the analytical need for a cause and the emotional need for the cause to be manageable. The two channels, the thinking that constructs the explanation and the feeling that receives it as adequate, have produced a coherent experience without either one checking the other. And so the explanation stands, and the student carries it, and the institution moves on.

Studies of teacher attributions consistently show that educators explain the same behavior differently depending on the student's race, class, and perceived background, and this finding has been documented so many times in so many contexts that its persistence in practice represents a failure of will rather than a failure of knowledge.^3^ When a white student from an affluent family acts out, teachers are more likely to attribute it to situational factors: something going on at home, boredom, insufficient challenge. When a Black student from a low-income family exhibits identical behavior, teachers are more likely to reach for dispositional explanations: disrespect, lack of self-control, inadequate upbringing. Okonofua and Eberhardt's (2015) research documented this with precision, showing that teachers' reactions to a second infraction were significantly more severe for Black students than for white students committing the same offense, and that this escalation was driven by the perception that the Black student's behavior reflected a "pattern" while the white student's behavior was an isolated incident.^4^ Gilliam, Maupin, Reyes, Accavitti, and Shic (2016) extended this to preschool settings, showing that when teachers were told to watch for challenging behavior, they looked at Black children, particularly Black boys, more than other children, before any behavior had occurred.^5^ The attribution process

began before there was anything to attribute. The groove was already selecting.

This is how the attribution groove becomes a mechanism for reproducing the conditions it claims to observe. When you believe Black boys are oppositional, you watch them more closely, interpret their behavior through a harsher lens, escalate consequences faster, and generate a disciplinary record that then becomes the evidence for the attribution you started with. The circularity is invisible from inside the process because each individual decision feels justified by the evidence available at the moment of the decision. The teacher who suspended Devon wasn't thinking about racial bias; the teacher was thinking about Devon's behavior in that particular moment, and the behavior was real, and the response felt proportionate, and only the pattern across hundreds of such decisions reveals what no single decision makes visible. The groove doesn't announce itself. It accumulates.

I find myself frustrated writing this, and I want to name the frustration rather than smooth it over, because the research has been clear for decades and the pattern has barely shifted. We know that deficit thinking drives disproportionate outcomes. We know that attributions shaped by race and class produce discipline disparities, identification disparities, tracking disparities, and expectation disparities. We know that the problem is structural rather than individual, that no amount of individual bias training addresses the systems and incentive structures that reward deficit attributions and punish institutional self-examination. And yet in meeting after meeting, year after year, the first explanation offered for struggling students locates the struggle inside the student, and the explanation settles into the room with the weight of obvious truth, and no one asks what else might also be true, because the groove has already provided an answer that satisfies both the analytical and the emotional needs of the people in the room. The students pay for that satisfaction. They have been paying for it for generations.

There is a specific attribution pattern that deserves examination because it is so pervasive in schools that it functions as the default operating language of educational failure. I call it the "they" move, and once you hear it, you cannot unhear it:

"They don't value education." "They can't sit still." "They won't do their homework." "They aren't motivated." "They need to take responsibility."

Each of these statements makes an attribution about the student or family or community, locates the cause of struggle in the characteristics of the person struggling, and accomplishes the same rhetorical work: absolving the institution. The "they" move is so automatic that educators often don't notice they're doing it; it feels like describing what's true rather than constructing an explanation that serves a function. But description and attribution are different operations, and the attribution has consequences. When you believe students "don't value education," you design interventions to change student values. When you recognize that students might be responding rationally to an education that has failed them and their communities for generations, you design

interventions that change what education offers. When you believe students "can't sit still," you build compliance systems. When you recognize that stillness is a culturally specific expectation that may not serve all learners, you build environments that work for different bodies and brains.

The attribution shapes the response, and the response either addresses the actual cause or doesn't. When it doesn't, the pattern persists, and the educator interprets the persistence as further evidence for their original attribution: "See? We tried interventions, and they still don't care." The loop closes, the groove deepens, and the student is now carrying both the original struggle and the institutional explanation of that struggle as a personal deficit.

There is a related move worth naming: the attribution of intent. When something goes wrong in a school, leaders often ask "Why did they do that?" as though the outcome resulted from a conscious choice. A teacher who refers students of color to the office at higher rates must be biased. A parent who misses conferences must not care. A student who won't complete assignments must be defiant. The question assumes intentional action and invites explanations rooted in character or motivation, but many outcomes result from systems, incentives, and constraints that shape behavior without anyone choosing them consciously. The teacher with the highest referral rates might also be the teacher with the most overcrowded classroom, the fewest support resources, and the least planning time, stressed beyond the capacity for the relational work that prevents escalation. Attributing to character what should be attributed to structure is one of the groove's most reliable tricks, because character explanations are simpler, more satisfying, and more available to conscious processing than structural explanations, which require holding multiple interacting causes in mind simultaneously, a cognitive task the brain reliably avoids when a simpler option is available.

I want to name the attribution pattern I find most dangerous, because it wears the clothing of equity analysis while functioning identically to deficit thinking.

Equity-conscious educators learn to see systemic factors. They understand that segregation, disinvestment, racism, poverty, and trauma shape educational outcomes. They know that blaming students for system failures is deficit thinking, and they're right about all of this. But sometimes the systemic analysis becomes its own attribution groove. I've watched educators explain student failure by pointing to systemic racism in a way that sounds critical but functions as absolution: "Of course these students struggle. Look at what systems have done to them. They've been traumatized. They've been deprived of opportunity. They've been subjected to chronic stress." Every word might be true, and yet the explanation locates the cause of failure outside the school's control and suggests the school should do very little, since the cause is societal. This is deficit thinking in progressive clothing: the content changed from "these kids can't succeed because of their culture" to "these kids can't succeed because of systemic oppression," but the function didn't change at all. Both explanations protect the institution from having to examine itself.

Real systemic analysis doesn't stop at naming the system. It asks: given these systemic conditions, what can this school do differently? It recognizes that students shaped by oppressive systems are not determined by those systems. It holds both truths, that systems matter and that schools have agency within systems, simultaneously, which is exactly what the groove resists, because the groove wants a single cause located in a single place, preferably outside the explainer's control. The attribution groove can wear any ideological outfit: it can sound conservative, blaming students and families, or progressive, blaming systems in a way that absolves the school. The tell is always the same. Does the explanation invite institutional change, or does it protect the institution from having to change?

Three disciplines interrupt the attribution groove. I practice all three inconsistently, and they work better on some days than others, which is its own kind of evidence about the depth of the grooves.

The first discipline is asking the alternative question: "What if the cause is something I can control?" When educators explain a student's struggle as family disruption, the alternative question is: "What if something within our school is contributing, and what would we look for?" This question doesn't deny external factors; it ensures you're considering internal ones. Most of the time, causation is multiple, with family factors, school factors, developmental factors, and community factors interacting in ways that simple attribution cannot capture. The groove selects for the explanation that protects the institution, and overriding it requires forcing yourself to examine institutional factors even when external explanations are available. External attributions feel safe; they're comfortable and don't require self-examination. Internal attributions are threatening because they imply that you might need to change. The groove protects against that discomfort, and overriding it takes deliberate effort that the rhythm of a meeting rarely affords.

The second discipline is seeking the success case. Whatever pattern you're trying to explain, somewhere in your system there's a counter-example: students who should be struggling but aren't, classrooms where the trend is reversed, schools that beat the demographic predictions. These cases don't prove your general analysis is wrong, but they complicate the attribution. If family disruption causes school failure, why did the student turn around when a scheduling accident placed him with a teacher who saw him differently? His family didn't change. Something else did. The success case forces you out of the groove because it presents evidence your groove can't explain, and accounting for it usually reveals something about what's possible that the dominant explanation obscures. I've started asking this question systematically: whenever someone offers a deficit-based explanation for student outcomes, I ask who's succeeding despite those same conditions and what we can learn from them. The question shifts attention from what's wrong with students to what's right with certain practices, and it opens space for institutional learning that the attribution groove forecloses.

The third discipline is the who-benefits audit. Every attribution has consequences: when we explain student failure as family dysfunction, families bear the consequence; when we explain teacher struggle as lack

CHAPTER 4

# The Prediction Groove

We have known about prediction failure for decades. The cognitive science is not ambiguous, not emerging, not contested in any serious way by anyone who has read the literature with care. Tetlock tracked 28,000 predictions made by political experts across multiple domains and found that their accuracy barely exceeded what you would get from a dart-throwing chimpanzee; Kahneman and Tversky documented the planning fallacy across dozens of contexts, from infrastructure projects to military campaigns to kitchen renovations, and found that humans consistently underestimate timelines, costs, and risks while overestimating benefits, and that this pattern persists even among experts who have watched identical projects fail in identical ways; Meehl demonstrated in the 1950s that simple statistical models outperform clinical judgment in prediction after prediction.^1^ We have researched it, documented it, published it in journals, taught it in graduate programs, cited it in keynote speeches, and built entire consulting practices around the insight that humans are bad at predicting the future. And we keep acting as if our predictions are facts, constructing futures in our heads, populating those futures with people who behave exactly the way we expect, and then treating the whole construction as the most likely outcome, as though the vividness of our imagination constitutes evidence for its accuracy. The prediction groove costs students time they do not have, and we keep making the same prediction errors with the same confidence while using the same language about being "data-driven" and "evidence-based," when what we are actually doing is building elaborate narratives about a future we cannot see and then defending those narratives against any evidence that they are wrong.

This is the third groove: the deep cognitive channel through which we construct the future before we have tested it, then treat our construction as reality. Unlike the evidence groove, which distorts what we see now, and the attribution groove, which distorts who we blame, the prediction groove distorts what we expect will happen. It is perhaps the most insidious of the three because it operates invisibly, because we do not experience our predictions as predictions but as reasonable assessments of probable outcomes, because the feeling of predicting and the feeling of knowing are, from the inside, nearly indistinguishable. That indistinguishability is the mechanism: the body registers a constructed future with the same felt certainty it registers a remembered past, and because the feeling is the same, the mind treats both as equally real. The prediction is not experienced as speculation. It is experienced as knowledge, carried in the body with the authority of something that has already happened, and the analytical process that should be testing the prediction instead defends it, because the felt certainty arrived first and shaped what the analysis would find.

The prediction groove manifests in several distinct patterns, all sharing a common feature: false certainty about events that have not occurred.

Fortune-telling is the act of predicting specific outcomes with a confidence that reality does not warrant. A leader preparing for a difficult conversation does not suspect the other person will become defensive; she knows it, with a felt certainty that masquerades as insight, that registers in the body as knowledge rather than speculation, that carries the authority of experience even when the experience being drawn upon comes from different people in different contexts about different topics. The leader who has decided that her superintendent will reject the equity proposal has not weighed probabilities; she has written a script, cast the roles, rehearsed the dialogue in the shower and the car and during meetings where she should have been paying attention to other things, and now the script feels like a memory of something that has already happened rather than a guess about something that has not. The sensation in the body of the predicted future and the somatic signal of a recalled past are processed by the same cognitive architecture, which is why predictions acquire the solidity of memory and why challenging someone's prediction feels, to them, like challenging their experience.

Catastrophizing is the specific form of fortune-telling that constructs worst-case futures with minimal probability adjustment: the conversation will not just be uncomfortable but will destroy the relationship; the initiative will not just face resistance but will fail spectacularly and take the leader's credibility with it, and everyone will remember she was the one who pushed for it, and her career in this district will never recover. Notice how the catastrophic chain extends itself, each link adding emotional weight that the analytical process mistakes for evidence. Catastrophizing does not merely predict negative outcomes; it predicts maximally negative outcomes and assigns those outcomes a probability approaching certainty, because the vividness of the imagined catastrophe feels like evidence for its likelihood, because the emotional weight of the worst case drowns out the cognitive work of assigning realistic probabilities to multiple possible outcomes. The thinking channel is constructing the catastrophe while the feeling channel is responding to it as though it has already occurred, and the two channels are reinforcing

each other in a feedback loop that produces the conviction that the worst case is not merely possible but imminent.

Status quo bias operates through prediction by constructing change as inherently risky while framing the current state as inherently stable, even when the current state is doing measurable harm to identifiable students whose names we know and whose cumulative files we have reviewed.^2^ Leaders overestimate the probability that new initiatives will fail and underestimate the probability that current practices are already failing, which means that inaction feels like the conservative choice when it is often the most radical choice available, the choice with the most predictable consequences for the students who can least afford to wait while adults calculate their risk tolerance.

Optimism bias works in the opposite direction, constructing futures where our initiatives succeed, our timelines hold, and our assumptions prove correct. Leaders underestimate the probability of obstacles and overestimate their own capacity to navigate them, and this is not the same as hope, because hope acknowledges uncertainty while optimism bias denies it; hope says "this might work, and I am willing to try," while optimism bias says "this will work because I have planned it, because I have read the research, because I am the kind of leader who gets things done." Both status quo bias and optimism bias are prediction errors, one making us too cautious and the other too confident, and both feel, from the inside, like clear-eyed realism, which is precisely what makes them so difficult to interrupt.

Here is what I find most frustrating about the prediction groove in educational leadership: it forecloses testing. It prevents the very experiments that might reveal our predictions are wrong, and because we never test, we never learn, and because we never learn, our predictions feel validated by the outcomes we create by not testing them. A leader who has decided that a difficult conversation will go badly does not have the conversation, which means she never discovers that the other person has been waiting for months for someone to raise the issue, has seen the same data, has been constructing her own catastrophic predictions about what would happen if she raised it herself. Two leaders, both predicting the same failed conversation, both avoiding it for the same reasons, both paying the same cost in months of inaction while the students whose data prompted the concern keep getting sent out of classrooms. The prediction groove does not just distort individual judgment; it creates coordinated paralysis, because every leader in the system is constructing the same catastrophic futures about the same conversations, and no one is testing any of them. The felt certainty that the conversation will fail is shared, and because it is shared, it feels even more like knowledge, because "everyone knows" functions as confirmation even when everyone is predicting rather than knowing.

A planning committee can spend eight months designing a school for a community whose voices they predicted rather than solicited, because the committee chair has decided, based on experience with other communities in other contexts about other topics, that families will not engage meaningfully with questions

about pedagogy and school culture. Another constructed future treated as fact. When someone finally asks how she knows what families want, the pause that follows reveals the whole architecture of assumption: she does not know; she is predicting based on patterns she has seen elsewhere, and she has not tested whether those patterns hold for this community, with this kind of engagement, about this school. When the committee finally designs genuine engagement, through home visits conducted in multiple languages by trusted community members, they learn that their predictions were partially right and profoundly wrong: yes, families cared about practical matters like transportation and food, but they also had deep concerns about cultural responsiveness, about whether their children would see themselves reflected in curriculum and staff, about discipline approaches that diverged sharply from the committee's assumptions, about academic challenge that the "student-centered" vision had not emphasized. Eight months of design work shifted because eight months of predictions proved inaccurate, and the cost was borne by the families who were talked about rather than talked with.

The prediction groove feeds the evidence groove, and both grow stronger in the feeding. If a leader never has the difficult conversation, she never discovers her colleague agrees with her, and her prediction stands unchallenged; she might even gather what feels like confirmatory evidence afterward, a dismissive tone in an unrelated meeting, a budget decision that seems to deprioritize equity work, and each piece of evidence confirms what she already knows. The future she constructed becomes the only future she can see, and alternative futures, the ones where the conversation goes well, where the colleague has been waiting for it, where the relationship actually strengthens through honest engagement, become invisible, present in the field of possibility but absent from the field of felt experience, which is the only field the prediction groove recognizes.

The planning fallacy deserves particular scrutiny because it is the prediction groove operating at the institutional scale, shaping not just individual conversations but entire theories of change. Kahneman and Tversky identified it with precision: our systematic tendency to underestimate the time, costs, and risks of planned actions while overestimating their benefits.^3^ It is not occasional, not the result of insufficient information or inexperience or poor leadership. It is a predictable feature of how humans construct future projects, persisting among experts who have watched similar projects fail, who have published articles about why projects fail, who have consulted on failing projects and charged substantial fees to explain why the failure was predictable. Buehler, Griffin, and Ross extended this work to show that the planning fallacy holds even when people are explicitly warned about it, even when they are shown base rates for similar projects, even when they themselves have experienced identical failures in their own recent history.^4^ Lovallo and Kahneman later distinguished between the "inside view," where planners imagine themselves executing the project successfully, and the "outside view," where statistical base rates from comparable projects provide a far more accurate forecast, and found that even when leaders have access to both views, they overwhelmingly

default to the inside view, because the inside view is vivid and personal and feels like planning, while the outside view is statistical and impersonal and feels like pessimism.^5^

Consider how district initiatives typically unfold, because the pattern is so consistent it could be scripted. A superintendent announces a new literacy curriculum. The implementation timeline shows pilot classrooms in fall, full rollout in spring, measurable gains on state assessments by year two. The budget allocates funds for materials and initial training but not for the follow-up coaching that the implementation science literature has demonstrated is necessary for transfer to practice.^6^ The theory of action connects curriculum to instruction to learning with clean arrows and confident predictions, and the board approves it unanimously, because no one asks what happened with the last three curriculum adoptions, because asking would be awkward and because everyone in the room has already committed to the optimistic timeline.

Two years later, the assessment gains have not materialized. The timeline slipped by six months because materials arrived late and the pilot revealed problems no one anticipated, although the implementation science literature anticipated exactly those problems and described them in detail that was available to anyone who looked. The budget overran by forty percent because the training required more follow-up than anyone predicted, although the research on professional development transfer could have told them what to expect. The pilot classrooms revealed problems that required curriculum modifications, which delayed the rollout, which compressed the training, which compromised implementation fidelity, which undermined the outcomes the initiative was designed to achieve, which led to the conclusion that the curriculum was the problem rather than the implementation, which led to a new curriculum adoption two years later, which started the cycle again with the same confident predictions and the same planning fallacy. I have watched this cycle repeat in district after district, and what makes it so maddening is not that leaders fail to plan well but that they fail to learn from the last time they planned well and watched the plan unravel in exactly the ways the research predicted it would.

There is a version of the prediction groove that operates in service of change rather than against it, and I want to name it because it is particularly seductive for justice-oriented leaders. When you have done the analysis, read the research, and seen the patterns clearly, the urgency of the problem generates false certainty about the solution. A principal who sees clearly that compliance-oriented, worksheet-driven instruction is failing Black and Brown students may develop what I will call prophetic certainty: the conviction that she knows not just what needs to change but how change must unfold, on what timeline, against what resistance, and with what outcome. She moves fast because children are being harmed now, and the moral urgency is real, and the diagnosis is correct, and she has seen transformation work in other schools, and she knows the way forward.

Within months, she announces comprehensive changes: project-based learning across all grade levels, elimination of worksheets and packets, student-led conferences, radical revision of the schedule. The changes start immediately because waiting feels unconscionable. She predicts resistance from veteran teachers, and she

is right. She predicts some families will be confused, and she is right. She predicts the transition will be rocky but ultimately successful, and she is wrong, profoundly wrong, wrong in ways that cost real people real things. Her best teachers leave, exhausted by the pace of change and demoralized by what they experience as contempt for their existing expertise, for the relationships they have built with families, for the skills she never bothered to assess because her prophetic certainty did not require an inventory of what already worked. Family complaints escalate to the board. Assessment scores drop. Her successor quietly restores much of what she eliminated.

The prediction groove does not care whether your values are correct. It distorts the future regardless of whether you are defending the status quo or trying to dismantle it. Diagnosis is not prophecy, and seeing the problem clearly and predicting the solution accurately are entirely different cognitive operations. Conflating them is one of the most expensive errors a justice-oriented leader can make, because the students who pay for the error are the same students the leader was trying to serve. The interval between what she felt about the urgency, which was accurate, and what she predicted about the solution, which was wrong, is the prediction groove's signature: felt certainty about the problem transferring, without inspection, into felt certainty about the fix. One channel, the one registering the moral urgency, ran ahead of the other, the one that should have been testing assumptions about implementation, and the running ahead felt like clarity rather than separation.

The cost of prediction failure is not abstract. It is counted in the months that students spend being disproportionately disciplined while two leaders avoid a conversation they have each decided will fail. It is counted in school designs that do not reflect community voice because a committee predicted what families would say instead of asking them. It is counted in teachers who leave because a principal's prophetic certainty left no room for their professional knowledge. It is counted in the cumulative weight of initiatives that begin with confident predictions, encounter predictable obstacles, and collapse into the cynicism that makes the next initiative harder to launch, harder to staff, harder to fund, harder to believe in. The people who pay for our prediction errors are almost never the people who made them.

The core question for interrupting the prediction groove is: "What am I assuming?" Predictions depend on assumptions we rarely surface. When a leader predicts that a colleague will become defensive, she is assuming that past behavior in other contexts about other topics predicts future behavior in this context about this topic, that her colleague's public persona reflects private thinking, that the colleague has not already noticed the same patterns and wished someone would raise them. None of these assumptions are visible until you ask the question, and once they are visible, their fragility becomes apparent, and the constructed future that felt so solid begins to dissolve into what it always was: a guess dressed up as knowledge.

But the prediction groove does not respond to a single question the way a lock responds to a key. The groove has been carved by repetition, deepened by confirmation, and maintained by the fact that we almost

CHAPTER 5

# See What You're Missing

## The Evidence Discipline

What would it actually take to change your mind?

Not in the abstract, not the version of the question where you nod thoughtfully and say something about being open to new evidence. I mean concretely, specifically, on a decision you're currently defending: what evidence, arriving in what form, from what source, would cause you to abandon your current position and adopt a different one?

I ask because I'm not sure I can answer it myself, and I've been thinking about it for years. I have practices I'm about to describe in this chapter, practices for interrupting the evidence groove, for forcing different information into view, for catching confirmation bias before it distorts decisions that affect students. I believe in these practices. I've used them. I've watched other leaders use them. And I genuinely don't know whether they work reliably enough to justify the confidence with which I'm about to present them.

That uncertainty is the real content of this chapter, not a throat-clearing exercise before I get to the prescriptions. The evidence groove doesn't politely step aside when you learn counter-practices; it adapts, operates on the counter-practices themselves, shapes how you implement them, what you notice when you use them, and whether you interpret their results honestly. Every tool I'm about to offer you is vulnerable to the same distortion it's designed to correct. I think they help. I think they help more than doing nothing. But I've been wrong about what helps before, and the times I was most wrong were often the times I felt most certain, which is itself a data point about the relationship between certainty and accuracy that I want you to hold while you read what follows.

So here is what I can offer: a set of practices that seem to create friction against the groove's natural direction, offered with less confidence than I would like, because honesty requires it.

## The Counter-Practice, Such As It Is

Part One of this book named the problem. The evidence groove shapes what data we collect, how we interpret what we find, and whose voices count as credible. We see what confirms our existing beliefs and discount what challenges them, and we do this automatically, unconsciously, with complete confidence that we're being objective.

Knowing this doesn't fix it. You can understand confirmation bias perfectly and still fall into it every day, can teach workshops on selective attention and still notice only the evidence that supports your preferred conclusions, can read this entire book and walk into your next meeting seeing exactly what you expect to see. Knowledge is necessary but insufficient. What's needed is practice: repeated behaviors that create friction against the groove's natural direction, structures built in advance that force different information into view. At least, that's the theory. The theory is clean. The execution is messy, inconsistent, and shaped by the very biases it's meant to interrupt. I want to be honest about that messiness even as I describe the structures, because pretending the counter-practices work like clockwork would itself be a kind of evidence groove, one where I only notice the times the practices succeed and forget the times they don't.

The core discipline is simple to state and difficult to maintain: actively seek evidence that would prove you wrong. This is not natural. Our minds resist it. We experience disconfirming evidence as a threat, and we respond to threats by defending ourselves, and the defense happens in the body before it happens in the analysis. There's a felt contraction when someone presents evidence that challenges a position you've committed to, a tightening that the thinking channel then rationalizes as skepticism about the evidence rather than discomfort about being wrong. The instinct to protect our existing beliefs is ancient, deep, and powerful. But leaders who build this discipline discover something that I think is true, though I can't prove it with the certainty I'd like: being proven wrong early is much less painful than being proven wrong after you've committed resources, relationships, and reputation to a flawed course of action.

## The "What Would Change My Mind?" Protocol

Before committing to any significant decision, ask yourself: What evidence would convince me I'm wrong?

Most of us, when asked, discover we can't answer it. We know what evidence would confirm our view; we can describe in detail the data that would prove us right. But evidence that would change our mind is harder, sometimes impossible. Sometimes we discover that no evidence could change our mind, which means we're not holding a conclusion at all but a faith, and when you can't articulate what would change your mind, you're not making a decision but defending a position. Decisions can be revised based on evidence. Positions can only be abandoned or maintained, and we rarely abandon them willingly.

Consider what happened when Principal David Chen tried to address discipline problems at his school. He was convinced the problem stemmed from inconsistent enforcement across classrooms: some teachers were strict, others lenient, and students were confused about expectations. The solution, obviously, was standardization: common expectations, common consequences, common documentation. He built the system, trained the staff, implemented the protocols with fidelity. Six months later, referrals hadn't decreased; they had increased.

A colleague asked him a version of the question I opened this chapter with: "Before you implemented this, what would have convinced you it was the wrong approach?" David realized he had never asked. He had collected data on inconsistency because he was looking for inconsistency, had found it because that's what he was measuring, and had never specified what outcomes would signal that his theory was wrong. When he went back to the data with fresh eyes, he found something he had missed: the referrals weren't randomly distributed but concentrated in a handful of classrooms, during a handful of periods, involving students who were also struggling in specific academic areas. The problem wasn't inconsistent enforcement; it was a mismatch between academic challenge and student support in particular contexts. Standardizing consequences addressed a problem that wasn't the actual problem.

I tell this story as if it illustrates a clean lesson, and in a sense it does, but I should note that I've seen leaders ask the "what would change my mind?" question and still find ways to dismiss the disconfirming evidence when it arrives. The question is a starting place. You write down what failure would look like, commit to looking for it, and then face the much harder task of actually letting the evidence change your thinking instead of explaining it away. I'm not sure this protocol works as often as I want it to. I'm fairly sure it works more often than not asking the question at all.

## The Data Diet Audit

What information do you consume on a regular basis, and what information is systematically absent from your intake?

Your data diet shapes what you know, which shapes what you think is possible. Consider a typical principal's information diet: standardized test scores, attendance data, discipline referrals, parent complaints that escalate to the office, teacher requests for meetings, emails from central office, performance evaluations of staff they directly supervise. This is significant data, but notice what's missing. They don't routinely see how students experience the school moment by moment unless students complain loudly enough, don't see how teachers experience leadership decisions unless teachers push back forcefully enough, don't see what happens in classrooms when no administrator is present, which is most of the instructional day. Community perceptions remain invisible unless those perceptions generate visible conflict, and early warning signs of staff burnout stay hidden until those signs become resignations or medical leave.

The data diet creates blind spots that follow patterns. Information flows uphill in hierarchies, which means the principal naturally sees what subordinates push upward: problems get escalated while successes often don't, complaints reach the office while quiet satisfaction stays in classrooms. The data diet is biased toward whatever generates enough noise to break through.

A data diet audit asks: what types of information regularly reach me, what types don't unless I actively seek them, whose perspectives are represented and whose are absent, what decisions am I making based on incomplete information because the complete information never arrives at my desk?

Superintendent Rachel Okonkwo found that she regularly received information from principals, cabinet members, and parents who escalated complaints, but rarely from teachers, students, support staff, or parents who were satisfied or who had given up on being heard. She restructured her information flow: monthly listening sessions with random groups of teachers, student surveys that went directly to her office, channels for support staff and custodians and cafeteria workers to raise concerns without going through supervisors. She didn't abandon her existing data sources; she supplemented them with sources her position naturally excluded.

Did it work? Partly, I think. She heard things she wouldn't otherwise have heard. Whether she was able to act on what she heard with the consistency she intended, whether the new channels eventually calcified into their own grooves, whether the listening sessions became performative over time: these are questions I can't answer with confidence, and I notice my uncertainty about them as a kind of data about the limits of what structural fixes can accomplish when the minds operating the structures carry the same grooves the structures are designed to interrupt.

## Adversarial Input Structures

The evidence groove is strongest when everyone in the room agrees with you. Consensus feels good, feels like validation, like the group has done good work and reached the right conclusion together. But consensus can also mean the groove is operating across your entire team, everyone seeing the same confirming evidence and dismissing the same disconfirming signals, with the group's confidence increasing precisely as its accuracy decreases.

Adversarial input structures deliberately introduce friction into decision-making, forcing consideration of perspectives that don't emerge naturally from the group's shared assumptions. I want to describe three of them, with the caveat that each can be co-opted by the very dynamics it's designed to disrupt.

The devil's advocate role assigns someone to argue against the emerging consensus, identifying weaknesses, raising objections, and surfacing counterarguments whether or not they personally believe them. The role works best when it rotates regularly so no one becomes the designated negative presence, and when it's genuinely protected so the advocate doesn't suffer social consequences for effective criticism. But even well-implemented devil's advocate roles sometimes devolve into a kind of theater where the designated critic raises manageable objections that feel challenging but don't actually threaten the leader's preferred direction, and the group gets the satisfaction of having "considered the other side" without having genuinely risked changing course.

The pre-mortem goes further. Before implementing a decision, the team imagines it's twelve months in the future and the initiative has failed spectacularly. Working backward, each person independently writes down what went wrong: what caused the failure, what warning signs were ignored, what assumptions proved false. Research by psychologist Gary Klein suggests that pre-mortems increase the ability to identify reasons for potential failure by roughly 30%.^1^ The technique works, I think, because it shifts the psychological frame; instead of defending a plan we've already endorsed, we're explaining why it failed, which activates different cognitive processes and surfaces concerns that optimism bias would otherwise hide.

The red team is more intensive: a separate group, not involved in the original decision, tasked with attacking the plan, looking for flaws, weaknesses, and failure modes the planning team missed. They ask: if we wanted this to fail, how would we do it? If an adversary wanted to undermine this initiative, what vulnerabilities would they exploit? Red teams are common in military and security contexts but rare in education, which I think is a mistake, though I acknowledge that the cultures are different enough that transplanting the practice directly may not work as cleanly as I'd like to believe.

## The Disconfirmation Habit

The practices described above are structural: they can be implemented as organizational procedures, embedded in meeting agendas, required before certain categories of decisions. But the deepest counter to the evidence

groove, if there is one, is a personal habit: the discipline of actively seeking disconfirmation.

This means treating your current beliefs as hypotheses rather than conclusions, looking for the evidence that would challenge them because you want to be accurate rather than because you want to be wrong. It means experiencing confirmation of your beliefs with mild suspicion rather than satisfaction: if you keep finding evidence that supports your view, you might be right, or you might not be looking hard enough for evidence that challenges it, and the two are difficult to distinguish from the inside. There's a felt quality to each, a satisfaction that accompanies confirmation and a discomfort that accompanies challenge, and the evidence groove operates partly by steering you toward the satisfaction and away from the discomfort, below the level of conscious choice.

Superintendent Terrence Boyd built this habit through a practice I find compelling and also wonder about. Every time he felt certain about something, he asked himself who disagreed and why, not to change his mind automatically but to understand the strongest case against his position. When he was certain a school needed new leadership, he sought out people who believed the current leader could succeed. When he was certain a program wasn't working, he talked to its most passionate advocates. When he was certain about budget priorities, he explored the case for different priorities with people who would make different choices.

Terrence often remained convinced he was right after hearing the opposition, but sometimes he discovered he was wrong in ways that mattered. His certainty became informed rather than reflexive, and he could explain not just why he believed what he believed but why intelligent people disagreed. I find this practice admirable. I also notice that it depends on a kind of intellectual honesty that the evidence groove itself can compromise, and I don't have a clean answer to the question of how you know you're genuinely engaging with the opposition rather than collecting enough counter-evidence to feel rigorous while remaining fundamentally unchanged.

Dr. Amara Okafor kept a notebook, not a leadership journal or a reflection log, but a record of times she had been wrong. Every Friday afternoon, she spent fifteen minutes reviewing her week and asking a single question: what did I believe on Monday that turned out not to be true? She wrote down the belief, the evidence that challenged it, and what she learned. After six months, she had documented forty-seven instances of being wrong about something that mattered: wrong about a teacher she thought was disengaged who was actually grieving, wrong about a parent she read as hostile who was actually afraid, wrong about a curriculum intervention she was certain would work.

The notebook changed how Amara led, she told me, because it made her more cautious about her own certainty. She had documented proof that her first impressions were unreliable, evidence that her confidence didn't predict her accuracy. When you document your errors, you accumulate evidence about your own reliability as a judge, discover patterns in when and how you're wrong, and become appropriately humble because the data supports it rather than because humility sounds good in speeches.

I've recommended this practice to other leaders. Some have found it transformative. Others stopped after a few weeks because it was depressing, or because they couldn't identify their errors in retrospect, or because the whole exercise felt like self-criticism that didn't actually change their behavior. I don't know what to make of those failures. Maybe the practice doesn't work for everyone. Maybe those leaders were experiencing the evidence groove even in their attempt to counter it, unable to see the errors that a different perspective would have made visible. The groove operates on the counter-practice just as it operates on everything else.

## Voice Mapping and the Problem of Absent Perspectives

Every decision is shaped by the voices present when it's made, and the evidence groove operates partly through voice: we treat information from some sources as credible and dismiss information from other sources as unreliable, biased, or uninformed, and this follows patterns of power, proximity, and perceived expertise.

Voice mapping makes these patterns visible. Before a significant decision, ask: who has been heard in this process, who hasn't, who is in the room when decisions are made, who is absent, whose perspectives have shaped the options we're considering?

Curriculum Director Marcus Washington conducted a voice map before selecting a new math curriculum. The process seemed robust: committee meetings, teacher surveys, pilot data, vendor presentations. But the voice map revealed that the committee was composed primarily of representatives from high-performing schools, that teacher survey response rates varied dramatically by building with the lowest responses from schools with the highest poverty rates, and that pilot data came from classrooms led by early-adopter teachers. The voices shaping the decision were systematically unrepresentative.

Marcus restructured the process, adding focus groups in underrepresented schools, surveying students directly, holding community sessions in neighborhoods that didn't typically engage with district decisions. The additional voices complicated the decision, surfaced concerns the committee hadn't considered, and also surfaced assets and community knowledge the committee hadn't recognized.

There is a deeper problem here that I keep circling and never quite resolving. The people most affected by decisions are often the hardest to reach because they don't come to meetings, don't respond to surveys, don't have the time or the trust or the belief that their voice matters. You can build every structure I've described in this chapter and still miss the perspectives that matter most, because those perspectives belong to people who have learned through long experience that speaking up changes nothing. The evidence that reaches leaders is filtered by power structures, and people with more power have more access to leaders' attention, which means your evidence base systematically overrepresents the views of the powerful and underrepresents the views of the marginalized. Interrupting this pattern requires deliberate effort to invert the hierarchy, to seek evidence from those with least access, to weight their perspectives more heavily precisely because they're usually

underweighted. I believe this is true, and I also notice how easily this principle can become its own performance, a box checked rather than a practice sustained, the listening session held so the leader can say they listened rather than because they intend to be changed by what they hear.

## What Changes, and What Doesn't

Leaders who build these practices don't become perfectly accurate. The evidence groove doesn't disappear. You can do everything in this chapter and still miss critical information, still interpret data through biased frames, still discount voices that challenge your preferred narrative.

But something shifts. Certainty becomes qualified: instead of "I know this is the right approach," leaders begin saying "Based on what I currently see, this appears to be the right approach," and the qualification signals precision rather than weakness, a judgment based on evidence that could change. Resistance becomes information rather than threat: what is this person seeing that I'm not, what do they know that I don't? Decisions improve because the evidence available to the leader expands, the voices shaping their thinking multiply, the blind spots shrink enough to matter. And the gap between what the leader senses about a situation, the felt quality of something being off, and what the leader's analysis says about that situation narrows, because the practices create space for the felt signal to reach the analytical process before the analysis has already constructed its conclusion.

I've described these practices with more confidence than I feel. Every counter-practice I've offered is vulnerable to the evidence groove: the pre-mortem can become a ritual that surfaces predictable concerns rather than genuine surprises, the devil's advocate can learn to raise objections that feel challenging but don't threaten the preferred direction, the data diet audit can become a quarterly exercise producing a report nobody reads, the disconfirmation habit can become intellectual performance. Even Amara's notebook can become a catalog of small, safe errors that protects the leader from confronting the large, uncomfortable ones.

I don't know how to solve this. I don't know if it can be solved. The evidence groove is a condition of having a human mind, and the counter-practices are themselves products of human minds, which means they carry the same vulnerabilities they're designed to address. What I can say, with whatever credibility remains after all these caveats, is that the practices seem to help more than their absence. Imperfect friction against the groove is better than no friction at all.

But I want to leave you with the question I started with, because I still haven't answered it for myself. You've read five chapters now about how the evidence groove operates, how it shapes what you see and what you miss, how it bends your thinking toward confirmation and away from surprise. You've encountered practices designed to interrupt it, and you've heard me admit that those practices are themselves compromised by the very distortion they target.

CHAPTER 6

# Blame the System, Not Too Much

A high school principal reviews suspension data in her office. Black males are suspended at three times the rate of white males. She writes "implicit bias" on the whiteboard and circles it twice. The intervention she designs targets teacher mindset: unconscious bias training for all staff, mandatory by October. She does not examine the referral process, the code of conduct language, or the hallway supervision protocols that funnel certain students toward certain outcomes. The cause has been named. The naming feels like progress. It is not progress. It is the point where analysis stopped being useful.

A superintendent sits in a budget meeting. Three schools in his district are underperforming by every available metric. He attributes the pattern to chronic underfunding and says so publicly, repeatedly, with conviction. The underfunding is real. It is also twenty years old. During those twenty years, two schools with identical funding levels performed adequately. He does not mention them. The systemic explanation is true and it is incomplete, and the incompleteness protects him from examining what his central office does with the money it has.

A curriculum director presents to the school board. Reading scores are flat across the district for the third consecutive year. She attributes the stagnation to pandemic learning loss, a cause so large and so external that no one in the room bears responsibility for it. Board members nod. The explanation accounts for the pattern without implicating the curriculum adoption process she led, the professional development calendar she designed, or the pacing guides her office wrote. Pandemic learning loss is real. It is also a shield.

Three leaders, three settings, three attributions locating cause somewhere comfortable. The pattern is precision failure: each identified a real contributing factor and then stopped looking, because the factor they found protected them or their organization from harder questions. The cause they named was accurate. The place they stopped was convenient. And how far apart what each leader sensed about the inadequacy of their explanation, the felt quality that something was still missing, and what their analysis delivered as a finished answer, represents the attribution groove operating at the point where systemic thinking becomes its own form of avoidance.

This chapter is about that stopping point.

## The Precision Problem

Systemic thinking is a corrective. The attribution groove described in the previous chapter pulls toward individual explanations: this teacher cannot manage a classroom, this student lacks motivation, this parent does not value education. The corrective says look at structures, examine what the system makes easy and what it makes hard, consider what would happen if different people occupied the same positions. If outcomes would be similar regardless of who fills the role, the system is the cause.

The corrective is necessary. It is also dangerous when it becomes its own groove. Once a leader learns to see systems, systems appear everywhere, and every problem becomes a systems problem, every failure attributed to structures and policies and resources and histories. The analysis grows more sophisticated while the outcomes remain unchanged, because systemic analysis without precise targeting does not lead to action but to more analysis. The leader produces increasingly elegant explanations of what is broken while nothing gets fixed, and the elegance itself becomes a performance substituting for the harder work of locating the exact mechanism that needs to change. There is a felt satisfaction in naming the system, a sense of having done intellectual work that the body registers as adequate, and that satisfaction can become its own endpoint, the thinking channel running at full resolution while the action channel sits idle, waiting for an analysis that never quite finishes because the analysis has become the work rather than the preparation for it.

The discipline: blame the system, but be precise about which part. Not so broadly that the blame diffuses into abstraction. Not so comfortably that naming the system becomes a substitute for changing anything specific within it.

Accurate attribution identifies which elements of which systems produce which outcomes. That distinction sounds academic. It is the difference between a district that spends three years analyzing its equity gaps and a district that identifies a specific evaluation instrument producing racially disparate eligibility determinations and replaces that instrument in six months.

## Where the Cause Actually Lives

A behavior pattern in a school gets attributed to classroom management deficits among teachers. That attribution directs intervention toward teacher training. But the pattern might live somewhere else entirely in the system: in a schedule that creates forty-minute transitions between classes during which no structured supervision exists; in a discipline code that defines disruption so broadly that ordinary adolescent behavior triggers formal consequences; in a referral form that takes ninety seconds to complete, making removal easier and faster than any de-escalation strategy the district has taught.

The attribution "classroom management deficit" is not necessarily wrong. It is frequently imprecise, naming a category of cause without locating the mechanism, and the mechanism is where intervention happens.

Consider two possible attributions for the same pattern. The first: "teachers need better training." The second: "the referral process makes removal faster than any alternative response." The first is a dispositional attribution wearing systemic clothing, still locating cause in individual capacity, in what teachers know or do not know. The second identifies a structural incentive that would produce similar behavior from almost anyone occupying that position. The intervention implied by each is completely different: the first leads to more professional development; the second leads to redesigning the referral process so that alternative responses are equally accessible, equally fast, equally supported by administrative follow-through.

Precision matters because resources are finite and because misattribution compounds. A district that spends its professional development budget on bias training when the actual mechanism is a flawed evaluation instrument has not addressed the problem. The money is spent. The outcomes persist. Then the next attribution cycle begins, usually blaming the people who received the training for not applying it with sufficient fidelity.

Disproportionality in special education provides a clear example. A cabinet-level analysis attributes the pattern to teacher bias in the referral process, directing intervention toward teacher mindset. But the actual data might show something different: referral rates roughly proportional, with disproportionality emerging at the evaluation stage where instruments with embedded cultural bias convert referrals into eligibilities at different rates for different populations, or at the eligibility determination stage where committee judgments are shaped by limited cultural knowledge and narrow developmental norms, or still further downstream in communication with families so inadequate that misunderstandings about student history distort how information gets interpreted.

Each location in the pipeline implies a different intervention. Referral-stage disproportionality calls for work with teachers. Evaluation-stage disproportionality calls for examining instruments. Eligibility-stage disproportionality calls for restructuring committee processes. Communication failures call for redesigning how families are engaged throughout. The attribution "bias causes disproportionality" is not wrong. It is too

imprecise to guide action toward the right part of the system.

## The Controllability Question

When an attribution settles on a cause, the next question is whether that cause can be changed by the people in the room. This sounds obvious. It is routinely skipped.

Poverty causes disparate outcomes. True. Not changeable by a school district. Historical disinvestment created the conditions being observed. True. Not reversible by a building principal. Systemic racism produces differential access to opportunity. True. Not solvable by a curriculum coordinator in her office.

These attributions are accurate and they are, functionally, explanations for why nothing can be done. If the cause is beyond organizational control, the attribution becomes an endpoint rather than a starting point, explaining the problem with satisfying completeness while producing zero intervention. The explanation is correct, and the students are still failing, and the correctness of the explanation does not help them.

The discipline is to accept external causes as real and then ask: given these causes, what within this organization's control affects the outcome? This is not denial of systemic forces; it is the refusal to let systemic analysis become an alibi for inaction.

A school cannot eliminate poverty. It can control what happens during the hours students are present: the quality of instruction, the responsiveness to early warning signs, the depth of relationships adults build with young people, the way families are communicated with about what matters and why. None of these eliminate the external cause. All of them represent points where the organization has genuine agency, where different choices produce different results for the students currently enrolled.

When attendance problems get attributed to transportation barriers, housing instability, and family work schedules, the attribution may be entirely accurate. The question is what follows. If the answer is "these are systemic issues beyond this school's control," the attribution has become a wall. If the answer is "given these realities, where in this school's response can different conditions be created," the attribution becomes a floor. Same-day contact for absences rather than waiting until three days have passed. Restructured morning arrival to reduce the cost of being late by fifteen minutes. Identification of which students have a strong relationship with at least one adult and which do not, followed by deliberate action to close that gap. Examination of which periods have the highest rates of students leaving early, followed by restructuring those periods so that what happens during them is worth staying for. None of these solve poverty. All of them represent choices within the school's control.

The controllability question prevents a specific failure mode: the leader who delivers increasingly sophisticated systemic analyses while outcomes for students remain unchanged year after year. The analysis is

correct. The students are still failing. The sophistication of the explanation provides no benefit to the people it explains.

## Learning from Deviation

When something goes wrong, the instinct is to ask why. When something goes right, the instinct is to feel relief and move on. The second instinct is more costly than the first.

Within any system producing poor outcomes, there are deviations: students who succeed despite identical risk profiles, classrooms with different results serving similar populations, buildings that outperform prediction. These deviations contain information that failure analysis misses, because failure analysis asks "why is this broken" while deviation analysis asks "why does this work here when it does not work there."

If chronic absenteeism sits at twenty-eight percent across a school, the natural analysis focuses on why students are absent, surfacing the usual causes: transportation, family circumstances, disengagement, health challenges. These causes are real and also uniform across the student body in ways that the absenteeism is not. Scattered through the data are students with identical risk profiles who attend consistently, classrooms where attendance runs ten points higher than the building average despite serving the same population, periods where students remain when they leave during other periods with different teachers or different content.

These deviations are evidence that the systemic explanation, while true, is incomplete. Something additional operates in those cases, enabling different outcomes within the same constraints. Identifying what that something is requires a different analytical posture: instead of "why is attendance bad," the question becomes "where is attendance better, and what conditions exist there that do not exist elsewhere in this same building." The answers tend to be specific and actionable rather than abstract and immovable: strong extracurricular connections giving students reason to be present beyond obligation, relational attendance practices where adults notice and name absence personally rather than processing it bureaucratically, engaging instruction during the periods students are most likely to leave.

None of these findings are surprising in retrospect. They were not visible until the analytical question changed from failure-focused to deviation-focused. The failure question produced explanations that were true and immovable. The deviation question produced explanations that were true and actionable.

## The Convenience Audit

Every explanation of a problem serves someone's interests. This is not a claim about bad faith; it is an observation about the consequences explanations carry beyond their accuracy.

An attribution that locates a problem in students protects teachers from scrutiny. An attribution locating it in teachers protects administrators. An attribution locating it in administrators protects the board. An attribution locating it in funding protects everyone inside the system while asking someone outside to act.

The audit: when an explanation gains acceptance, ask who is absolved, who continues without change, who avoids the discomfort of examining their own contribution.

A district explains flat math scores as a teacher fidelity problem: teachers are not implementing the new curriculum as designed, adapting and modifying and supplementing and deviating from the intended sequence. The explanation feels rigorous because it identifies a specific cause and implies a specific intervention: enforce fidelity. But look at who is protected. The curriculum selection process remains unquestioned. The training calendar, compressed into too few days with no follow-up mechanism, remains unquestioned. The pacing guide, written without reference to the actual prerequisite knowledge students possess, remains unquestioned. The fidelity metrics, measuring procedural compliance rather than the adaptive expertise effective teaching requires, remain unquestioned. The administrators who approved these decisions remain unquestioned. The only people not protected are the teachers absorbing the blame for failing to implement a curriculum they were inadequately prepared to teach, using materials that assumed resources their classrooms did not have, at a pace their students could not sustain.

The convenience audit does not mean that explanations protecting the powerful are always wrong. Sometimes teachers genuinely fail to implement, and sometimes the fidelity explanation is accurate. But when an explanation conveniently shields those with authority from examination while concentrating blame on those without authority, that pattern warrants investigation: not assumption of conspiracy, but investigation of whether the explanation is accurate or merely convenient.

## Holding Disposition and Situation Together

The attribution groove pulls toward dispositional explanations: this person has these characteristics, which produce these outcomes. The systemic corrective pulls toward situational explanations: this person is in these circumstances, which produce these outcomes. Each pull, taken alone, fails.

If everything is dispositional, every outcome is earned, every failure deserved, and the system bears no responsibility for the conditions it creates.

If everything is situational, no individual bears responsibility for anything, and changing behavior requires changing the entire architecture first, which changes slowly, and the people managing the slow change get to keep managing it indefinitely.

CHAPTER 7

# Test What You Predict

Dr. Nkechi Okafor sat at the head of the conference table on a Tuesday afternoon in October, watching her leadership team flip through the twelve-page strategic plan she had spent the better part of a year building, and she could feel the room doing what rooms always do when people are proud of something they made together: nodding at the goals, murmuring approval at the timelines, treating the projections as if they were weather forecasts issued by professionals who had instruments for measuring wind. The plan was, by every conventional standard, excellent. Three years of aligned reading and mathematics targets, graduation rate benchmarks, family engagement metrics, teacher retention projections, and a theory of action so internally coherent that the board had approved it unanimously and the community had endorsed it at two consecutive public meetings. Nkechi herself had staked a great deal on it, not just her professional reputation but her private belief that careful planning could tame the uncertainty that made leadership feel, on its worst days, like steering a ship by studying the wake.

"I think we're ready," said Marcus, her assistant superintendent, closing his binder with the satisfied snap of someone who has reached the end of a long process and wants to celebrate the reaching.

She nodded. They were ready.

None of the predictions came true.

Not the reading gains, not the mathematics improvement, not the graduation rate increase, not the family engagement metrics or the teacher retention numbers or the budget projections that assumed the state would maintain funding levels, not the technology implementation timeline that assumed vendor reliability that never materialized, not the professional development impact targets that assumed protected follow-up time evaporating the first week of November. Three years later, Nkechi stood before the same board that had approved the plan and explained why almost nothing had worked as predicted, and the explanations were numerous and each was individually true: an unexpected budget crisis in year two, leadership turnover in year three, vendor failures on critical systems, resistance she had not anticipated from sources she had considered allies. All true and all beside the point, because the plan had failed not due to bad luck or bad actors but because it was built entirely on untested predictions, every goal assuming interventions would work as designed, every timeline assuming obstacles could be navigated as planned, every projection assuming the future would unfold as the planners imagined, with no mechanism anywhere in the document to test these predictions early, catch them failing, and adjust while adjustment was still possible.

There was a moment, early in the implementation, when Nkechi had felt something was off. The first quarterly data showed attendance gains but no academic movement, and instead of reading that signal as information about the theory of action, her analytical process explained it as an implementation lag, a normal feature of first-year rollouts. The felt signal that the plan might be wrong and the analytical explanation for why the signal could be ignored were running on separate tracks, and the analytical track, with its reassuring language about implementation timelines and reasonable expectations, was louder. The separation between what she sensed and what she analyzed cost three years.

Nkechi had learned strategic planning from consultants who taught her to set ambitious goals and build coherent theories of action. What no one taught her was that plans are hypotheses rather than prophecies, and that they need to be tested rather than implemented with fidelity and defended against evidence of failure. This chapter is about building the discipline of testing: actively probing your predictions to discover where they are wrong before the wrongness becomes catastrophic and irreversible.

## Predictions Are Hypotheses

The prediction groove, as the previous chapter described, constructs futures and treats them as reality. We predict how conversations will go, how initiatives will unfold, how people will respond to change, and we experience these predictions as reasonable forecasts grounded in experience and analysis, feeling realistic when we are actually being imaginative in ways we cannot see from the inside.

The counter-practice begins with a reframe: treat every prediction as a hypothesis requiring evidence. A hypothesis is a tentative explanation that needs testing; it might be right, it might be wrong, and you do not know until you gather evidence. Scientists do not attach their identity to

hypotheses, because attachment is considered a liability rather than a virtue, a sign that the investigator has confused what they want to be true with what they have reason to believe is true.^1^

Leaders rarely think this way about their predictions. The strategic plan is treated as a commitment rather than a hypothesis. The initiative timeline is a promise to stakeholders rather than a guess. The expected outcomes are what the leader will be held accountable for delivering, which makes prediction testing feel like weakness rather than wisdom, like hedging rather than rigor. But the reframe is more honest than the alternative, because you do not know that the strategic plan will work; you believe it will, based on reasoning that may or may not be sound. You do not know the timeline is achievable; you have estimated it, using assumptions that may or may not hold. You do not know the outcomes will materialize; you are predicting they will, based on a theory of action that may or may not connect cause to effect in the real world as cleanly as it connects them on paper.

Treating predictions as hypotheses does not mean abandoning confidence or commitment; it means holding that confidence provisionally, subject to revision based on what you learn, and building mechanisms to test predictions early, when revision is still possible and the cost of being wrong has not yet been distributed across an entire system and the students it serves.

## The Pre-Mortem

Before launching any significant initiative, gather your team and ask them to imagine it is one year in the future and the initiative has failed completely. Not struggled, not underperformed: failed, with the board asking what went wrong and the community demanding accountability. Working independently, each person writes down what caused the failure: which assumptions proved false, which obstacles emerged that no one anticipated, which warning signs were ignored because the group was committed to success and could not hear what commitment did not want to hear.

Then share and compare. The failure modes that multiple people imagined independently are your highest-probability risks; the unique failure modes that only one person saw might be the surprises that blindside you precisely because they live in a single person's peripheral vision rather than the group's shared field of view.

The pre-mortem works because it shifts the psychological frame.^2^ When a group has decided to pursue something, criticism feels like threat, and people explain away concerns or dismiss the person raising problems as negative, as someone who does not believe in the work. The pre-mortem gives permission to name problems by making problem-identification the explicit task, converting what would otherwise feel like disloyalty into an act of professional responsibility the group has collectively sanctioned.

Principal Tomoko Hayashi introduced pre-mortems before every major initiative in her building, and the first few were uncomfortable in exactly the ways discomfort is supposed to be uncomfortable. People were not used to imagining failure before they had started, and some participants wrote superficial failure modes rather than engaging genuinely. But the practice caught problems that normal planning missed: a new teacher evaluation system, imagined through pre-mortem, revealed that it depended on training capacity that did not exist and principal buy-in that had not been built; a technology rollout revealed assumptions about bandwidth and device reliability that proved false when tested in pilot classrooms. The pre-mortem surfaced these gaps before implementation, when they could be addressed through additional planning rather than discovered through failures that damaged trust and wasted years.

The pre-mortem does not prevent failure. It surfaces the predictions most likely to be wrong, so you can test them before committing fully, monitor them during implementation, or build contingencies in advance rather than scrambling after the fact.

## Scenario Planning

The prediction groove constructs one future and lives in it; scenario planning constructs multiple futures and prepares for several. The technique emerged from military and business strategy, where the consequences of being wrong about the future are severe and sometimes fatal, and it asks three questions the prediction groove prefers to leave unasked: what are the key uncertainties that will shape how the future unfolds; for each uncertainty, what are the plausible alternatives; and for each combination of alternatives, what world might we find ourselves in?

Educational leaders face genuine uncertainties that analysis cannot resolve regardless of how sophisticated the analysis becomes: will state funding increase, stay flat, or decrease; will enrollment grow or shrink; will the political environment support innovation or demand tradition; will housing costs push families into or out of the district. These are not questions you can answer through better data but uncertainties you have to plan around, and the prediction groove's preference for a single confident forecast is precisely what makes leaders vulnerable when the forecast proves wrong.

Superintendent James Washington faced a bond measure that polling showed as uncertain, and the prediction groove would have picked one outcome, probably the optimistic one, and planned accordingly, leaving scrambling as the only option if the prediction proved wrong. Instead, his team built three futures: the bond passes comfortably, the bond fails narrowly, the bond fails significantly. For each future, they developed response strategies, early indicators that would signal which future was materializing, and transition plans that could be activated within days. When the bond failed narrowly, James was not scrambling; he had already imagined this future and prepared for it, and while the failure was still painful, it was not catastrophic because it

had been anticipated and planned for rather than experienced as an unimaginable shock.

Scenario planning takes more time than single-future planning and requires imagining outcomes you do not want, which is cognitively expensive and emotionally unpleasant. But it produces plans that survive contact with reality rather than shattering on first impact.

## Small Bets

The prediction groove scales up before testing: you are confident the initiative will work, so you implement it everywhere simultaneously; you are sure the approach is right, so you commit fully before seeing results. Small bets reverse this logic. Before full commitment, test the prediction at limited scale; before district-wide implementation, pilot in selected sites; before assuming the conversation will go one way, have a preliminary conversation to test your assumptions.

Curriculum Director Angela Park was convinced her new literacy intervention would transform outcomes, and the research supported it, and the vendor assured her with case studies, and the theory of action was coherent. Her confidence was high, but she had been confident before about other interventions that had not worked as predicted, and she had watched other curriculum directors implement with confidence and fail with surprise that might have been avoidable. So instead of district-wide rollout, she proposed a pilot: four schools, one semester, clear metrics for success before expansion.

The pilot revealed problems the research had not prepared her for. The intervention worked beautifully in schools with strong instructional coaches but floundered in schools without that support infrastructure, worked with teachers who had flexible approaches to curriculum but created conflict with teachers who preferred more structure, worked at certain grade levels better than others for reasons that only became clear through implementation and that no amount of pre-implementation planning could have surfaced. None of this invalidated the intervention; it refined the prediction. The intervention worked, but only under conditions that had not been part of the research summaries or vendor presentations, and full rollout would need to either ensure those conditions or target schools where conditions already existed.

The pilot cost one semester in four schools. Full rollout without piloting would have cost years and millions, with failures distributed across the entire system and affecting thousands of students who did not have years to wait for adults to learn what a pilot could have taught them in months.

## Reference Class Forecasting

When you predict how your initiative will unfold, you are taking an inside view: imagining your specific plan, your specific team, your specific context, and constructing the future from the details you know intimately.^3^ Reference class forecasting takes an outside view. It asks a question the prediction groove does not want asked: when others attempted similar things, what actually happened? Not in their imagination before implementation but in reality after, not in their optimistic projections but in their actual results measured by people who had no stake in the answer.

The technique: identify the reference class, meaning similar initiatives in similar contexts at similar scale, and look at the distribution of actual outcomes. How long did they take? How much did they cost? What percentage achieved their stated goals? The results are often sobering. Curriculum implementations typically take twice as long and cost fifty percent more than planned. School construction projects average thirty percent over budget and eighteen months behind schedule. Professional development initiatives show measurable classroom impact less than twenty percent of the time.

These statistics feel abstract until you apply them to your own plans. Your curriculum implementation is budgeted for eighteen months; the reference class says thirty-six months is more likely. Your construction project is budgeted at fifty million dollars; the reference class says sixty-five million is more realistic. Your professional development is designed to change practice; the reference class says there is an eighty percent chance it will not, at least not in measurable ways. Reference class forecasting does not mean you will match the average, and maybe your plan is better designed, your team stronger, your context more favorable. But if you are predicting outcomes significantly better than the reference class achieved, you need a specific reason why, and "because we're committed" is not a reason; it is the prediction groove talking.

I want to be honest about something that makes me uncomfortable. I have used reference class forecasting to revise timelines and budget projections, and I have watched it work, and I have also watched myself quietly set the reference class data aside when it told me something I did not want to hear about an initiative I was personally invested in. The discipline is real, but it is not automatic, and knowing about the prediction groove does not inoculate you against it. I still catch myself constructing optimistic futures and defending them with the same confidence I warn others about, which is either a humbling reminder that cognitive grooves run deeper than knowledge, or evidence that I should not be writing this chapter, or possibly both.

## Testing Interpersonal Predictions

Not all predictions are about initiatives and timelines. Many of the predictions shaping daily leadership are about people: how someone will respond to feedback, how a colleague will react to a proposal, what a community will say about a decision. These interpersonal predictions determine which conversations happen and which get avoided, which ideas get voiced and which die in the silence between what a leader thinks and

what they are willing to say out loud.

Interpersonal predictions are particularly prone to the groove because they feel grounded in knowledge of the people involved. You know Dr. Carter, so you can predict how she will respond to criticism, and this felt-knowledge makes interpersonal predictions feel more like facts than guesses, more like reading the room than constructing a story about it. But interpersonal predictions are often wrong. People surprise us when we give them the chance, and small changes in how we approach a conversation can produce outcomes that our predictions never imagined because our predictions were based on how past conversations went rather than how this conversation might go under different conditions.

The discipline is to test interpersonal predictions rather than acting on them. Before avoiding the conversation you have predicted will go badly, have a smaller conversation that tests your prediction. Before assuming someone will resist your proposal, ask them what they think before you have committed to a position that makes their response feel like either agreement or attack.

Assistant Principal Ravi Mehta had predicted for months that his principal would not support a schedule change he wanted to propose. He knew her: traditional in her approach, valued stability, did not like disruption to established routines. So he did not bring it up, building quiet resentment as the schedule continued to create problems he could see solutions for but could not implement without her support. When he finally tested the prediction by mentioning the schedule problem in passing, her response surprised him. She had noticed the same problems. She had been thinking about changes. She wondered why no one on her team had raised it, because she assumed they were satisfied with the current arrangement.

His prediction had kept him silent about an issue his principal was ready to address. Two people, both sensing the same problem in their bodies, both constructing analytical explanations for why the other person would not want to discuss it, both paying the cost of silence while the students whose schedules needed changing waited for adults to test a prediction that was wrong. What lies between what Ravi sensed about the schedule and what his prediction said about his principal's receptiveness was the prediction groove's signature: felt knowledge about a person substituting for actual knowledge, constructed certainty preventing the conversation that would have dissolved it.

## What Changes

Leaders who build prediction-testing disciplines do not become better at predicting. The future remains unknowable, and the prediction groove still operates, still constructs futures, still generates false confidence about outcomes that have not happened. What changes is the relationship to predictions: they become provisional rather than fixed, hypotheses rather than facts, things to be tested rather than defended.

CHAPTER 8

# The Real-Time Protocol

The meeting was fourteen minutes in when I realized I had already decided.

Someone was presenting a proposal to restructure the master schedule, making room for more intervention time. The data was solid. The logic was clear. And somewhere in my head, a voice was already drafting the objections I would raise, not because the proposal was bad but because I hadn't thought of it, because the current schedule bore my fingerprints, and a better proposal meant I had missed something three years earlier when I designed it.

I caught myself. Not after the meeting, not that night in the shower replaying what went wrong, but in the moment. Fourteen minutes in. Something in my body registered before my thinking did: a tightening in the jaw, a readiness to speak that had nothing to do with what the presenter was actually saying. The objections I was composing had the quality of a reflex rather than a response, and I recognized the quality because I had felt it before, many times, without recognizing it. This time I wrote three words in my notebook: *What am I missing?*

Then I listened differently. Not for confirmation of my objections. For what the presenter might be seeing that I wasn't.

It took about thirty seconds. That's all the SEE Protocol requires when you've practiced it. Thirty seconds to interrupt the groove before it determines your response.

I want to tell you that those thirty seconds changed everything, that the protocol worked perfectly, that I became a different leader in that meeting. I'm not sure any of that is true. What I can say is that I asked different questions than I'd planned to ask, and the conversation went somewhere it wouldn't have gone otherwise. Whether that makes the protocol reliable or whether I just got lucky on a Tuesday afternoon, I honestly don't know.

There is another meeting I should tell you about, because this one did not go the way the first one did. A budget reallocation discussion, six people around the table, and I felt the groove activate early: a tightening in my certainty about which programs should be cut. I asked the question. "What am I not seeing?" A colleague answered honestly. She named a community partnership that would collapse if the line item disappeared, a partnership I had never considered because it did not appear on the spreadsheet I was reading. The answer was useful. It was accurate. It was the kind of disconfirming evidence the protocol is designed to surface. And the meeting continued as if she had not spoken, because the superintendent had already signaled a preference, and the political momentum in the room was carrying the decision forward regardless of what the question revealed. I watched the groove win. I had caught it. I had named it. I had surfaced the missing evidence. And the organization absorbed the catch the way a river absorbs a thrown stone: a brief disturbance, then the current resumes. I sat with the realization, afterward, that catching the distortion and changing the outcome are two different things, and the protocol only does the first.

This chapter is about using SEE in real time, not as a retrospective tool for analyzing past failures but as a live practice for catching distortion while there's still time to choose differently.

The three questions are simple enough to hold in memory:

*What am I not seeing?* (The evidence question.) *Who am I blaming?* (The attribution question.) *What am I assuming?* (The prediction question.)

Simple to remember, harder to ask when it matters. The challenge isn't intellectual, because the questions aren't complicated. The challenge is that your grooves don't want to be interrupted; they feel like clarity, like experience speaking, like the wisdom you've accumulated over years of doing this work, and questioning them feels like questioning yourself. Which it is. That's exactly why it's uncomfortable, and exactly why I'm not confident that telling people to "just ask the questions" is sufficient advice.

I've taught this protocol in workshops. I've watched leaders nod along, write the questions on index cards, stick them inside their notebooks. And I've watched many of those same leaders, six months later, describe situations where they forgot to ask, or asked but didn't listen to the answer, or listened but couldn't bring themselves to act on what they heard. I've done all three myself. The protocol is simple. The human being

using it is not.

The thirty-second version works like this, at least in theory.

You notice a signal, something in your body or your thinking or your emotional state. The signals differ for everyone. For me, it's a particular kind of certainty, a felt sense that I already know how this conversation is going to go, a settling in the chest that feels like confidence but is actually the groove clicking into its track. Other people describe different signals: composing a response before the other person has finished speaking, a tightening in the chest when someone challenges a decision, a flash of annoyance that masks the discomfort of having assumptions questioned.

The signal is your early warning system. It tells you the groove is activating, and the signal lives in the body before it reaches the analysis, which is why catching it requires a different kind of attention than the analytical attention most leaders have been trained to use. The analytical channel is already building its case by the time you notice the signal, and the signal is not a thought; it is a sensation, a quality of readiness or tightness or certainty that the thinking will rationalize if you give it enough time. The practice is to catch the sensation before the rationalization begins.

When you notice the signal, you pause. An internal pause, a slight slowing of your mental machinery. In that pause, you run the questions, not all three every time. You learn which question interrupts your dominant groove. For me, it's usually the evidence question: *What am I not seeing?* For leaders whose groove runs toward attribution, the question might be: *Who am I blaming?* For those trapped in prediction loops: *What am I assuming will happen?*

The pause creates space. In that space, you can choose. Do I proceed with my initial response, or does the question reveal something I need to attend to first?

Thirty seconds, sometimes less.

I describe this as if it's clean and sequential, as if the signal arrives and you calmly pause and select your question and proceed with new clarity. It doesn't feel like that. It feels rushed and uncertain and frequently incomplete. You catch the signal but the pause is too short. You ask the question but the meeting moves on before you can sit with the answer. You notice the groove activating and you interrupt it and then five minutes later you realize you've slipped right back into it without noticing. The thirty-second version is less like a protocol and more like a reflex you're trying to build, one that works sometimes and fails other times and improves so gradually that you can't always tell whether you're getting better or just getting better at telling yourself you're getting better.

The deep dive is different, and I have more confidence in it, though not as much as I'd like.

It's for decisions that matter enough to warrant deliberate analysis: hiring, program changes, resource allocation, responses to conflict. For the deep dive, you work through each question systematically.

Evidence: What data am I relying on? Where did it come from? Who collected it, and what questions shaped its collection? What data have I not sought? Who would have access to information that challenges my current view? Have I talked to them? If not, why not?

Attribution: What am I explaining, and what causes am I assigning? Am I locating the problem in individuals or in systems? Am I locating it in the people with the least power? What would the people most affected say is causing this? What explanation would be most uncomfortable for me to accept, and what would require me to change something about myself or my leadership?

Prediction: What am I assuming will happen? What evidence supports that assumption? What would falsify it? Have I made this prediction explicit enough that I could be proven wrong? Am I predicting based on hope or based on tested patterns?

The deep dive takes longer, fifteen minutes or sometimes more. I want to say it changes outcomes, and I've seen situations where it appears to, but I've also seen leaders go through every question with apparent seriousness and still arrive at the conclusion they started with. The questions can become ritual; you can learn to ask them in ways that feel rigorous but don't actually threaten your existing thinking. I've done this myself, gone through the full protocol and felt good about my thoroughness and later realized I'd been performing examination rather than doing it. The distance between genuine self-examination and the performance of self-examination is invisible from inside the performance, which is the same structural problem the grooves themselves create: the feeling of having done the work substitutes for the actual work, and the substitution is undetectable by the person making it.

Let me tell you about a curriculum adoption process, a composite of several I've been part of or heard about from colleagues.

A committee had been meeting for months. The data was in, the pilots had run, they were ready to recommend, and the leader had a strong sense of which curriculum should be selected.

She paused. "Before we finalize, I want to run a quick check. What are we not seeing?"

The room was quiet. People had already made up their minds.

Then a teacher spoke. "We haven't talked to families. We've piloted with teachers. We've looked at student data. But we haven't asked families what they think about homework expectations, or how they want to support reading at home, or whether they can access the online components."

The leader felt a flicker of annoyance, because this would delay the decision and they were so close. She noted the annoyance. That was the signal: the body registering a threat to the conclusion the analytical process had already reached.

"What would we learn if we asked?" she said.

What they learned shifted the recommendation. The curriculum they'd been leaning toward had excellent teacher supports but assumed family engagement patterns that didn't match their community: thirty minutes of homework support each night, devices many families didn't have reliable access to, parents attending daytime reading workshops. The second-choice curriculum had weaker professional development but better alignment with how families actually functioned, including audio support for parents who weren't confident readers, flexible homework expectations, and family components that didn't assume middle-class schedules.

The evidence question, asked in real time, surfaced information that months of analysis had missed, not because the committee was careless but because their grooves had defined "stakeholders" in a way that excluded the families whose children would be most affected.

I tell this story in workshops, and it lands well. People nod. They see the point. What I don't tell them, because I'm not sure how to say it without undermining the lesson, is that I've also seen the evidence question asked and answered and then ignored anyway. I've seen leaders surface the missing perspective and then explain why the timeline doesn't allow for it, why the budget has already been allocated, why the decision needs to be made by Friday. The question works when the person asking it is genuinely willing to be changed by the answer. I don't know how to teach that willingness. I'm not sure it can be taught.

I want to tell a different kind of story here, one I've been reluctant to include because it complicates the lesson, and because I don't come out of it looking as capable as the curriculum story does. A budget meeting, a proposal to cut a program I had built, and somewhere in the first few minutes I noticed the groove activating: the jaw tightening, the mental machinery starting to compose objections that had more to do with authorship than analysis. I wrote the question in my notebook. *What am I not seeing?* I sat with it. An answer came: the program was expensive to sustain relative to the number of students it served deeply, and the resources it consumed might serve more students differently. That was a real answer. It arrived clearly. I acknowledged it to myself and then watched the meeting continue past it anyway, because three other people in the room had already aligned on the cut before I got there, and the political momentum was carrying the decision forward regardless of what the question had revealed. The protocol worked. The groove was interrupted. I saw something I hadn't been seeing. The outcome was unchanged. I sat afterward with the realization that catching the groove and changing what happens are two different things, and the protocol only does the first. That distinction doesn't appear on the index card.

The structural embed is when you build SEE into your processes so thoroughly that the questions get asked even when you forget to ask them. I find this more promising than individual practice, though it has its own failure modes.

Every significant proposal includes a section titled "What This Analysis Doesn't Show," forcing presenters to name the limitations of their own data before the committee reviews it. Every hiring debrief begins with someone assigned to argue for the candidate not selected, not as theater but as genuine exploration of what might have been missed. Every initiative running more than six months gets a prediction audit: someone pulls out the original predictions and compares them to actual outcomes, asking where we were right, where wrong, and what the gap reveals about how we were thinking.

For any decision affecting a particular group, someone on the team is assigned to represent that group's likely perspective: before deciding on a new attendance policy, someone thinks from the perspective of families dealing with unstable housing; before changing the grading policy, someone considers the perspective of students working jobs after school. It's not the same as actually asking those people. The shadow perspective is a guess carrying the grooves of the person doing the guessing. But it's better than never considering the perspective at all, and it sometimes raises questions that lead to actually going and asking.

The structural embed appeals to me because it doesn't depend on individual discipline, which I've found unreliable in myself and in the leaders I work with. The questions get asked because the process demands it. But structures can become rituals too: the "What This Analysis Doesn't Show" section becomes boilerplate that everyone fills in without thinking; the prediction audit becomes a performance that everyone endures without learning from. I've seen both happen.

The most important moment for SEE is the moment when you feel most certain.

This is counterintuitive. We tend to question ourselves when we're uncertain and trust ourselves when we're confident, but the grooves are most powerful when they produce certainty. The brilliant blunder feels like insight. I've learned, slowly and imperfectly, to treat my own certainty as a signal rather than a conclusion. When I find myself thinking "obviously," that's when I try to write it down, because my "obviously" is often where my thinking is most distorted, where my grooves have constructed a reality that feels complete precisely because they've excluded the information that would complicate it.

I keep a running list, at least when I remember to. At the end of some months I review what I was certain about. Some of those certainties turn out to be correct, but a surprising number, maybe one in four, turn out to be wrong in ways that would have been predictable if I'd asked the questions. The pattern I've noticed: my strongest certainties cluster around what I want to be true. I'm rarely as certain about things I'd rather not face.

This is the paradox of clear thinking, or at least the paradox as I understand it: the moments when you most trust your judgment are often the moments when your judgment is least trustworthy, because your grooves are producing the sense of clarity that protects them from scrutiny, because the felt certainty in the body and the analytical confidence in the thinking are running together so smoothly that the smoothness itself becomes invisible, and the invisibility is the trap.

Back to that meeting, fourteen minutes in.

What I wasn't seeing was my own investment in the existing schedule. A better proposal meant I had missed something, and the objections I was drafting weren't about the proposal; they were about protecting my ego. The evidence question revealed an attribution problem: I was preparing to find flaws in someone else's thinking because I didn't want to face flaws in my own.

I didn't share any of this in the meeting. I simply asked different questions than I'd planned to ask. Instead of "Have you considered the complications with passing time?" I asked "What problems does this solve that the current schedule doesn't?" Instead of looking for weaknesses, I listened for strengths.

We adopted a modified version of the new schedule. It's been better for students. Intervention attendance is up. Teacher feedback is positive. The complications I was ready to raise turned out to be manageable. And I want to credit the protocol for that outcome, but I'm aware that I'm telling you this story because it worked, not telling you about the times I caught the groove and asked the questions and still made the wrong call, or the times I didn't catch it at all. The stories we choose to tell about our own growth are themselves shaped by the grooves we're trying to interrupt.

I want to be honest about what I don't know.

I don't know if the protocol works the way I think it does. I've seen it help. I've seen it fail. I've used it myself and found it useful on some days and performative on others, and I can't always tell the difference in the moment. The thirty-second version might be genuinely interrupting my grooves, or it might be giving me a sense of rigor that substitutes for actual rigor. The deep dive might be surfacing real blind spots, or it might be teaching me to ask better questions while leaving my deepest assumptions untouched. I don't have controlled studies. I don't have outcome data. I have my own experience, and my experience is exactly the kind of evidence this book has taught you to distrust.

What I can say is this: the protocol creates a pause. In that pause, sometimes, you see something you wouldn't have seen otherwise. The pause is where the two channels, the one carrying the felt signal and the one building the analytical case, have a chance to meet, and when they meet, the quality of the decision changes. What you do with it is still leadership, still judgment, still your responsibility. And the practice is never finished, because the grooves adapt. As you get better at catching one pattern, another emerges. You don't

CHAPTER 9

# Building Interruption Culture

Individual skill is insufficient. A leader who has mastered the SEE protocol, who can identify grooves in real time and interrupt them with discipline, still operates inside a system that shapes what questions get asked, what concerns get voiced, and what challenges reach the table. The research on organizational silence is unambiguous: the primary determinant of whether people speak up is not individual courage but organizational culture.^1^ When the environment punishes challenge, even skilled interrupters go quiet, and the grooves migrate from individual cognition to collective behavior, becoming harder to see precisely because everyone shares them.

This is the central problem of interruption as individual practice. One leader running SEE improves one leader's decisions, but a leadership team that has normalized interruption improves every decision the team touches. The cognitive distortions described in this book do not live only in individual minds; they live in groups, in shared assumptions that teams develop over years, in patterns of reasoning that become so familiar no one examines them. Irving Janis documented this phenomenon decades ago in his analysis of groupthink: cohesive groups develop shared cognitive frames that filter out disconfirming information and suppress dissent, not through explicit prohibition but through subtle social pressure that members often cannot identify.^2^ The pattern extends beyond the cognitive: a team's collective grooves produce a felt quality of agreement that registers in the body as consensus, and the bodily experience of consensus becomes its own evidence that the analysis is correct, so that questioning the conclusion feels like disrupting a state the group has worked to achieve rather than testing a hypothesis the group has not yet validated.

Chris Argyris called these "defensive routines": the organizational habits that prevent people from examining the assumptions most in need of examination.^3^ Defensive routines are self-sealing because they protect themselves from detection by making it socially costly to name them.

What makes this particularly dangerous in educational leadership is that the stakes are absorbed by people who are not in the room. When a leadership team's collective grooves distort a decision about curriculum adoption, resource allocation, or disciplinary policy, the consequences land on students and families who had no voice in the process and no mechanism for challenging its assumptions.

Building interruption culture is the organizational answer to an organizational problem.

## The Architecture of Organizational Silence

Before examining what interruption culture looks like, the pattern it must displace deserves scrutiny. Elizabeth Morrison and Frances Milliken's research on organizational silence reveals that most organizations develop systematic tendencies to suppress upward information flow.^4^ Employees at every level withhold concerns, disagreements, and negative information from those above them in the hierarchy, and the withholding is not random but concentrates around the topics that matter most: strategic direction, leadership effectiveness, ethical concerns.

The mechanism is straightforward. People watch what happens to those who speak up. If a colleague raises a concern and receives genuine engagement, the cost of speaking appears low. If that colleague receives dismissal dressed as appreciation, if the concern gets "tabled," if the relationship cools in the weeks that follow, every observer recalibrates. James Detert and Ethan Burris found that employees' willingness to voice concerns was predicted less by their individual disposition than by their assessment of leadership receptivity.^5^ People are remarkably accurate readers of whether challenge is genuinely welcome or merely tolerated.

This creates a cultural memory of challenge: every team carries a story about what happens to the person who disagrees, written not by announcements or mission statements but by accumulated observations of actual responses to actual challenges. The story shapes all subsequent behavior, often for years after the original incidents.

The implication is uncomfortable: if a leadership team operates in silence on important questions, that silence is almost certainly a product of leadership behavior, not necessarily current leadership since cultural memory can outlast the leaders who created it, but the silence is a response to something.

## Permission, Practice, and Protection

Building interruption culture requires three structural conditions: permission, practice, and protection. Each is necessary, none is sufficient alone, and the framework is sequential in theory but recursive in practice because each element reinforces or undermines the others continuously.

**Permission** means explicitly naming that challenge is expected rather than merely tolerated. Amy Edmondson's research on psychological safety demonstrates that teams perform better when members believe they can raise concerns without interpersonal risk.^6^ But Edmondson is careful to distinguish psychological safety from niceness; psychologically safe teams are not conflict-free but teams where disagreement about the work does not become a referendum on the person.

Permission requires specificity. "My door is always open" is not permission. "Before we finalize this decision, I need someone to articulate the strongest case against it" is permission. The specificity matters because vague invitations are easily read as performative, and in most organizations they are. Permission also requires repetition, because one bad response to challenge can undo months of permission-granting; a principal who responds to a teacher's concern with "I appreciate that, but we need to move forward" has just taught the entire staff that appreciation is the sound the door makes when it closes.

**Practice** means building structural mechanisms that make interruption routine rather than heroic. Without structure, challenge depends on individual courage, and courage is an unreliable organizational strategy. The pre-commitment question, the designated dissenter, the silence protocol all serve to routinize what would otherwise require someone to take a personal risk. Practice also means tolerating the discomfort of genuine disagreement; Argyris observed that most organizations have deeply ingrained norms against surfacing conflict, norms so embedded that participants experience conflict avoidance as politeness rather than suppression.^7^

**Protection** means ensuring that people who challenge do not suffer for it, either immediately or over time. The research on voice and silence consistently finds that the primary barrier to speaking up is fear of negative consequences.^8^ Protection addresses that fear through demonstrated behavior rather than reassurance. Most retaliation against challengers is invisible to the person doing it: the challenger's proposals face slightly more scrutiny than others, their name surfaces less often for visible projects, the warmth in the relationship cools by imperceptible degrees. Leaders who genuinely believe they welcome challenge often punish it through the same unconscious grooves this book has been examining, assigning dispositional labels to the challenger, filtering subsequent interactions through a lens shaped by the discomfort the challenge created. Protection therefore requires active monitoring: asking in private whether those who have challenged experienced any

cost, watching for patterns, and accepting that the answer may not be what the leader hoped.

## Structural Practices

**The Pre-Commitment Question.** Before presenting a decision, the team identifies in advance what evidence would change their minds. The responses are recorded and revisited after the decision plays out, creating retrospective accountability: if the team said "we would change direction if X happened," and X happens, the pre-commitment creates pressure to actually change direction rather than construct post-hoc rationalizations for staying the course.^9^

**The Designated Dissenter.** For consequential decisions, one team member is assigned to construct the strongest possible case against the proposed direction, not the casual invitation to "play devil's advocate" that produces token objections but a genuine mandate to find the objections the group has not considered. The role rotates and carries no social penalty because it is structural rather than personal.

**The Silence Protocol.** After a concern is raised, the team observes ten seconds of silence before any response. This interrupts the reflexive "but we already considered that" and "actually, the reason we didn't" that shut down information flow before it begins. The silence creates space for additional concerns to surface, because when people see that the first concern was not immediately dismissed, they recalibrate their assessment of risk.^10^

**The Steel Man Requirement.** Before critiquing a position, the critic must first articulate that position more strongly than the person who offered it: "Let me make sure I understand. You're saying that..." This prevents the pattern of arguing against caricatures rather than actual positions and slows the rush to defend, creating conditions for genuine understanding as a prerequisite for genuine disagreement.^11^

**The Anonymous Input Channel.** For high-stakes decisions, a mechanism for raising concerns without attribution. This supplements rather than replaces public challenge, particularly valuable in early stages of building interruption culture when the risks of public dissent remain high.

**The Outcome Review.** After major decisions have played out, the team revisits them to calibrate future decision-making: what was predicted, what occurred, where did the analysis fail. Tetlock's research on forecasting accuracy demonstrates that systematic outcome review is one of the few interventions that reliably improves subsequent prediction quality.^12^ The review should occur regardless of whether the decision succeeded, because success can be as instructive as failure when organizations succeed for reasons entirely different from what they predicted.

## Stages of Cultural Development

Cultural change follows a recognizable sequence.

**Stage One: Announcement.** The leader declares that challenge is welcome and introduces practices. Most efforts begin and die here. The announcement creates formal permission, but formal permission operates against the weight of cultural memory. People test the declaration cautiously, offering mild disagreements to see what happens. If the leader responds to even one test with defensiveness, deflection, or the appearance of appreciation without the substance of engagement, the effort stalls.^13^

**Stage Two: Testing.** Team members offer probes: small challenges with limited risk, a slightly critical question, a gentle alternative perspective. The leader's task is disciplined responsiveness to every probe: asking follow-up questions, acknowledging the value of the concern specifically rather than generically, and making visible changes based on what was raised. "Based on what you just identified, I'm changing my thinking about this" is worth more than any number of announcements about openness. This stage can last months, and attempting to accelerate it signals impatience that undermines the process.

**Stage Three: Emergence.** Someone offers genuine challenge with real stakes: a fundamental disagreement with the leader's direction, raised publicly. This is the moment that determines whether the culture advances or retreats. The leader's response does not need to be agreement; it needs to be genuine engagement, sustained attention, follow-up questions, visible wrestling with the implications, and either a changed decision or a transparent explanation of why the challenge was considered but not adopted. What the team is watching for is whether challenge is genuinely metabolized or merely acknowledged before being discarded.

**Stage Four: Normalization.** Challenge becomes ordinary. Concerns surface without drama. Disagreement happens without interpersonal tension. Edgar Schein's research on organizational culture emphasizes that cultural norms require continuous reinforcement; they are sustained by ongoing behavior rather than by historical precedent.^14^ A single leadership transition can reset a team to Stage One. A period of organizational stress can trigger regression. The work is never finished.

## The Authority Problem

Positional authority creates a specific and well-documented distortion in information flow.^15^ The more authority a leader holds, the less accurate the information that reaches them. People filter what they tell leaders, amplifying good news and softening bad news and omitting criticism that might damage the relationship. The filtering is usually unconscious on both sides.

For superintendents, this means operating with a more distorted picture of organizational reality than their principals have. For principals, it means receiving more distortion than their teachers. The hierarchy itself functions as a filter, and each level removes information that the level above needs.

Building interruption culture at higher levels of authority requires proportionally greater effort because the filtering is stronger. The leader must work harder to create safety, respond more carefully to challenge, and develop supplementary channels for accessing the information the hierarchy filters out.

There is a particular form of self-deception available to leaders with strong positional authority: the belief that they are exceptions to this pattern. The superintendent who says "my team is very honest with me" is almost certainly experiencing the filtering and interpreting it as honesty. The absence of challenge feels like agreement. The absence of criticism feels like approval. The feeling is wrong, but it is internally consistent, which makes it resistant to correction. The felt quality of agreement in the room, the sense that everyone sees the same thing, is itself the distortion operating at the collective level, producing the same false certainty in a group that the groove produces in an individual mind.

## The Objection of Efficiency

A predictable objection: this will slow everything down.

The objection reflects a miscalculation about where time is actually spent. Interruption culture takes more time in meetings and saves substantially more time afterward. The decision that gets challenged and refined before implementation does not generate the cascade of problems that unexamined decisions reliably produce. The concern surfaced early does not become the crisis consuming weeks of attention later.

Every leader has experienced the alternative: the fast decision that created slow problems, the initiative that launched without adequate challenge and then required months of remediation, the policy adopted on the strength of a compelling narrative that collapsed on contact with the conditions it failed to anticipate.

Argyris distinguished between single-loop learning, which corrects errors within existing assumptions, and double-loop learning, which examines the assumptions themselves.^16^ Interruption culture is the organizational infrastructure for double-loop learning. Without it, teams correct surface errors while leaving deeper distortions intact, guaranteeing that similar errors will recur.

The organizational conditions that determine whether interruption happens are not incidental to the work of clear thinking. They are the work. Individual cognitive discipline matters, but it operates within a system that either amplifies or suppresses it. A leader who has learned to interrupt their own grooves but works within a

CHAPTER 10

# When the Grooves Fight Back

Solomon Eze had done the work. Not the kind of work where someone attends a workshop and collects a certificate and goes back to the office feeling informed, but the real kind. He read an earlier draft of this book. He practiced the SEE Protocol for months before his district adopted it. He built structural embeds: prediction journals, pre-mortems, external review panels. He talked about interruption culture in cabinet meetings and modeled it in front of his board. He trained his principals. He revised his hiring protocols. He created a standing agenda item in every leadership team meeting that forced the group to surface disconfirming evidence before moving to a vote. He was, by every reasonable measure, a leader who understood cognitive grooves and had committed to fighting them.

He hired the wrong person anyway.

His district needed a chief academic officer. Solomon ran what looked like a rigorous process: diverse committee, structured interviews, rubric scoring, multiple rounds of deliberation. He asked his team to surface concerns. He invited challenge. He asked, in meetings, "What am I not seeing?" He did everything the framework says to do, and he did it with the fluency of someone who had internalized the language, who could explain to you why each step mattered, who could cite the research behind each structural embed.

The candidate he hired was articulate, confident, and deeply aligned with Solomon's vision. She referenced the same scholars Solomon cited. She used frameworks that mapped onto his own. She spoke about equity in language that resonated with how Solomon thought about equity. In other words: she was a mirror. The evidence groove operated inside the interruption protocol itself, and the protocol became the vehicle for confirmation rather than the check against it. Solomon was using the tools of clear thinking to build a more sophisticated version of the distortion the tools were designed to prevent.

Nine months later, the hire struggled. Not because she lacked competence, but because the match between her approach and Solomon's masked a critical gap: she had never led implementation in a district with Solomon's demographics, had never managed the specific kind of resistance his schools produced, had never been tested against conditions the interview process did not simulate. The concerns Solomon's team had raised during the process, concerns about implementation experience, about whether the candidate's equity language reflected depth or polish, had been raised and addressed and dismissed through the very protocol Solomon had built to prevent exactly this kind of dismissal. The questions were asked. The answers were accepted. The groove won.

I am frustrated by this story, and I want to be specific about why: it suggests that the practices described in this book, practices I believe in and have used and have recommended to hundreds of leaders, can be defeated by the very distortions they are designed to interrupt. The groove adapts. It learns the language of interruption and uses that language as cover. A leader who knows about confirmation bias can construct more elaborate confirmation architectures than one who doesn't, because the knowledge of the bias provides a false sense of immunity that the bias itself then exploits. Solomon wasn't ignorant of the grooves. He was sophisticated about them, and the sophistication became the newest layer of the trap.

This is what I mean when I say the grooves fight back. They do not accept disruption passively. They adapt, finding new channels when old channels are blocked, wearing new disguises when old disguises are recognized, and the more sophisticated the leader's understanding of cognitive distortion becomes, the more sophisticated the distortion's defenses become. The arms race is asymmetric: the leader must catch every groove every time, while the groove only needs to slip through once to determine a decision. And the groove has the advantage of operating below awareness, in the body, in the the body's reading of rightness that precedes and shapes the analytical process the leader believes is doing the thinking.

Solomon's experience illustrates the most dangerous adaptation: the groove that mimics the counter-practice. When a leader has internalized the language of interruption, "What am I not seeing?" and "Who am I blaming?" and "What am I assuming?" those questions can become performative rather than genuine, asked because the protocol requires them rather than because the leader is genuinely uncertain about the answer. The felt quality of the questioning changes: genuine uncertainty, the kind that lives in the body as discomfort and openness, gets replaced by procedural compliance, the kind that lives in the body as the satisfaction of having done the thorough thing. The replacement is invisible from

inside the experience, because both genuine uncertainty and procedural compliance feel like careful thinking. The difference is that one is actually checking the groove while the other is the groove running through the check without interruption.

The adaptation takes several recognizable forms.

The first is what I call the sophistication defense. As leaders develop more complex understandings of cognitive bias, their ability to construct elaborate justifications for their decisions increases proportionally. A leader who knows nothing about the evidence groove makes simple errors. A leader who understands the evidence groove thoroughly can construct an argument that accounts for the evidence groove, demonstrates awareness of confirmation bias, acknowledges the limitations of the analysis, and still arrives at a conclusion shaped by the groove, because the sophistication of the acknowledgment provides the bodily registration of rigor that substitutes for the actual rigor of changing one's mind. The more the leader knows about the trap, the more convincingly they can describe the trap while standing inside it. This is not hypocrisy. It is the groove adapting to the cognitive environment the leader has created. The leader built defenses. The groove found a way through them. The gap between what the leader feels about the adequacy of their analysis, which feels thorough, and what is actually happening in the analysis, which is confirmation dressed as self-awareness, is the groove's newest hiding place.

The second is the counter-practice fatigue. Leaders who maintain interruption practices for months eventually experience the practices as burdensome rather than essential. The pre-mortem becomes another meeting that takes too long. The prediction journal becomes another task on an overloaded list. The evidence question becomes a ritual repeated so frequently that it loses its investigative quality and becomes instead a formula, asked and answered with the same automatic cadence as "How are you?" in a hallway. The fatigue is genuine. Leadership is exhausting, and cognitive interruption adds load to an already overloaded system. But the fatigue is also the groove's strategy for reasserting itself: by making the counter-practice feel like a burden, it creates the conditions for abandonment, and abandonment returns the groove to its default position, operating without interruption, producing the felt certainty that the leader had learned to question and now no longer does.

The third is the identity capture. When a leader builds their professional identity around being someone who catches grooves, the identity itself becomes the newest groove. The leader now needs to be the person who sees clearly, who catches what others miss, who models interruption culture. That identity produces its own evidence groove: the leader selectively notices instances where the protocol worked and constructs a narrative of continuous improvement. It produces its own attribution groove: when errors occur, they are attributed to insufficient implementation rather than to limitations of the framework itself. It produces its own prediction groove: the leader predicts that continued practice will produce continued improvement, treating the trajectory

as guaranteed rather than contingent.

This is the deepest adaptation, and it mirrors a pattern the later books in this series examine in detail: the moment when the framework designed to liberate becomes the framework that constrains, when the tool for seeing becomes the lens that filters, when the sophistication of the self-examination becomes the newest obstacle to genuine self-knowledge. The leader who has built an identity around interrupting grooves cannot easily examine whether the interruption itself has become a groove, because the examination would threaten the identity, and the identity is what the leader uses to do the examining.

What does this mean practically? That the practices described in this book have a shelf life, that they need to be refreshed and challenged and revised, that a leader who has been using SEE for three years should expect the protocol to have calcified in ways that require detection and disruption, and that the detection requires something the protocol itself cannot provide: an outside perspective that is not embedded in the leader's cognitive architecture.

This is why the practices in this book are insufficient alone and why the series continues. The Logic Trap addresses the surface: the visible distortions in evidence, attribution, and prediction that you can learn to catch with practice. But beneath the surface, the patterns that shaped the grooves in the first place are still operating, shaping what you believe before you begin to think, shaping who you see yourself to be, shaping what your body registers as true before your mind has examined it. The deeper layers are the subject of the books that follow, and they matter because the grooves don't just filter information; they filter the counter-practices designed to filter the information. The arms race continues all the way down.

Solomon's story doesn't have a clean resolution. He didn't fire the hire and start over with perfect clarity. He adapted, slowly, painfully, over months: restructuring the role, providing support the original process hadn't anticipated, adjusting expectations, and examining what the failure revealed about the limits of his own interruption practices. The examination was productive and also incomplete, because every examination is, because the examining mind carries the same grooves it is trying to examine. The practice continues. The grooves persist. The span between the two narrows over time, which is the only honest metric of progress.

I am aware that this chapter risks being discouraging. You've spent nine chapters building understanding of the grooves and learning practices for interrupting them, and now I'm telling you the grooves adapt, the practices calcify, and the arms race has no end. If that feels like a reason to stop trying, I understand the impulse but reject the conclusion.

The alternative to imperfect practice is no practice. The alternative to a protocol that sometimes works is no protocol at all. The alternative to catching some distortions and missing others is missing all of them. The students whose futures depend on your decisions need a leader who is fighting the grooves, even imperfectly,

CHAPTER 11

# The Mirror and the Window

The framework presented in this book is itself vulnerable to the pattern it describes.

The SEE Protocol can become a tool for self-exemption: a sophisticated apparatus for analyzing other people's distortions that, precisely because of its sophistication, makes the user more confident in their own clarity. Kruger and Dunning noted that training people to recognize incompetence in others did not reliably improve their ability to recognize it in themselves.^1^ The same may be true here. Learning to identify grooves in others may not improve one's ability to identify grooves in oneself; it may make that identification harder by producing the illusion that one has already done the work.

Pronin, Lin, and Ross documented this as the bias blind spot: the consistent finding that people rate themselves as less susceptible to bias than others, and that this self-assessment becomes stronger, not weaker, when people are educated about bias.^2^ The more you know about cognitive distortion, the more confident you become that your knowledge protects you. It does not. The knowledge stays in the analytical channel. The bias operates in the channel underneath, the one that shapes which evidence reaches the analysis, the one that determines what the body registers as true before the mind begins its examination.

This chapter draws a distinction that determines whether the practices in this book produce genuine growth or sophisticated self-deception: the difference between the mirror, which examines your own thinking, and the window, which examines others' thinking. The mirror is primary. The window is derivative. Everything you see through the window is shaped by what you fail to see in the mirror.

The bias blind spot is the most reliable finding in the distortion literature.^3^ People believe they are less biased than others. They believe their judgments are more objective. They believe their self-assessment is accurate precisely because they have examined it, and the more thoroughly they believe they have examined it, the less receptive they become to external correction, because accepting correction would mean that the self-examination was inadequate, and accepting that the self-examination was inadequate threatens the competence the person uses to do the examining.

Ehrlinger, Gilovich, and Ross found that this effect intensifies with expertise.^4^ Experts are more susceptible to the bias blind spot than novices, because their expertise provides an additional layer of justification for their self-assessment. "I understand these biases. I've studied them. I teach workshops about them. Therefore I am less susceptible to them." The reasoning feels sound. It is the groove talking.

The practical consequence: the readers most at risk from this book are the readers who find it most persuasive, who internalize the framework most completely, who begin using the language most fluently, because the fluency itself creates a felt sense of immunity that the biases then exploit. If you've read this far and feel like you understand cognitive grooves well enough to avoid them, that feeling is the bias blind spot at work. The understanding stays in one channel. The groove operates in the other. And the felt certainty that understanding protects is itself the groove's newest product.

The mirror requires a specific discipline: turning the analytical tools in this book on your own thinking with the same relentlessness you would apply to anyone else's.

When you notice a colleague dismissing disconfirming evidence, you can identify the evidence groove operating in their thinking because you are outside their cognitive architecture. You have access to information they cannot see, and the seeing feels effortless because their grooves are not your grooves, their blind spots are not your blind spots, and the distance between observer and observed provides a clarity that self-examination cannot match.

The mirror removes that distance. When you examine your own thinking for the same patterns, the clarity dissolves because you are both the observer and the observed, both the examiner and the examined, and the grooves that shape your thinking also shape the examination of your thinking. You cannot step outside your own cognitive architecture to examine it from a neutral vantage point, because there is no neutral vantage point, because neutrality is itself a product of the architecture you are trying to examine. This circularity is not a problem to solve. It is a condition

to manage, and managing it requires accepting that self-examination is always partial, always compromised, always shaped by the grooves it is designed to detect, while still doing it because the alternative, no self-examination at all, is worse.

The preference for the window over the mirror is itself a groove, the groove of self-exemption operating in the precise domain where self-exemption is most dangerous: in the act of trying to help others think more clearly. The help one offers through the window is shaped entirely by the distortions one fails to catch in the mirror, and those distortions do not announce themselves. They feel like clarity, like insight, like the confidence that comes from seeing what someone else cannot see. That confidence is, more often than the confident person recognizes, the distortion itself.

Modeling is slower than correction. It is also more durable. A leader who catches their own grooves publicly, in meetings, in conversation, in the presence of subordinates, accomplishes several things without having to correct anyone: the possibility of catching grooves is demonstrated; the safety of acknowledging distortion is established; the grooves are normalized as features of all human thinking rather than symptoms of individual failure. Argyris observed that organizational learning depends less on what leaders say they value than on what leaders visibly practice.^5^ When a leader says "What am I missing here?" in a meeting, the question functions differently than when that same leader tells a colleague to ask "What am I missing?" The first is modeling. The second is instruction. People change their behavior in response to the first far more reliably than in response to the second.

When someone with authority demonstrates vulnerability about their own reasoning, they create permission for others to do the same. When someone with authority diagnoses others' reasoning, they create pressure to comply, which in most organizational contexts is the opposite of genuine reflection. A subordinate who performs the requested self-examination in order to satisfy a supervisor's expectation has not engaged in self-examination but in impression management, and the supervisor who mistakes one for the other has fallen into an evidence groove of their own.

Power dynamics structure every instance of this work. When a supervisor identifies a subordinate's groove, the subordinate faces a constrained response space where disagreeing with the boss carries real and often undiscussable costs. The "invitation" to see differently may function as a directive regardless of framing, and the subordinate may perform the requested reflection without actually engaging in it, producing surface compliance that looks like growth but changes nothing. The supervisor, receiving the performance, has no reliable method for distinguishing it from the real thing.

When a subordinate identifies a supervisor's groove, the risks reverse and intensify. The subordinate may face retaliation, the supervisor may dismiss the observation, or the supervisor may perform receptivity while

internally discarding the feedback. Each response is invisible to the subordinate, which means the subordinate receives no reliable signal about whether the intervention worked. The asymmetry is structural: the person with less power takes a greater risk and receives less information about its outcome.

When peers identify each other's grooves, the dynamics include competition, status, and relational risk. The observation may be heard as a power move rather than an act of collegiality, because in many organizational cultures, the ability to diagnose another person's thinking is itself a form of status, and offering an unsolicited analysis of a peer's cognitive distortion is a claim to superior insight.

None of these dynamics makes intervention inappropriate. All of them make it more complicated than the simple prescription "help others see their grooves" suggests. The words carry different weight depending on who speaks them and who receives them, and someone holding positional power who says "I might be wrong about this, tell me if you see it differently" is offering a genuine opening only if their history demonstrates that disagreement is actually safe.

There is a harder version of this mirror work, one that most readers will resist. Think of the last time someone offered feedback that felt unfair, that seemed to misunderstand the situation, that was dismissed as inaccurate, not feedback that was clearly wrong but feedback that produced irritation, the kind that gets explained away rather than absorbed.

What if they were right? Not entirely right, not right about everything, but what if the observation contained something accurate about a pattern easier to see from the outside than the inside? Pronin's bias blind spot research predicts exactly this resistance: the more confident a person is that they have examined their own reasoning thoroughly, the less likely they are to accept external correction, because the thoroughness of self-examination becomes a shield against information that contradicts it.^6^

If that question produces immediate defensiveness, if the first response is to re-litigate why the feedback was wrong, that response is itself data. The groove is defending itself, not necessarily because the feedback was correct but because the speed of the dismissal suggests it was never genuinely considered. The harder exercise is to return to that feedback now, weeks or months later, and sit with it long enough for the defensive architecture to become visible: the speed of the dismissal, the quality of the reasons marshaled against it, the certainty that arrived before any genuine consideration had occurred. That machinery is the groove operating in real time, and the only way to see it is to revisit the moment with enough distance for the defensiveness to lose its urgency.

The underlying questions retain their value independent of any framework. "What am I not seeing?" is useful with or without the evidence groove as a concept. "Who am I blaming?" prompts examination whether or not one has read attribution theory. "What am I assuming?" generates reflection regardless of theoretical

CHAPTER 12

# Staying in Practice

How long do you have to do this before it sticks?

I get asked some version of that question at nearly every workshop. Someone raises a hand, usually toward the end of the session, and asks it with genuine hope in their voice. They want a number: six months, a year, two years and then you're good. They want to hear that the grooves can be permanently rewired, that there's a finish line where the work of catching yourself becomes unnecessary because you've finally, fully learned.

I don't have that answer, and I'm not sure the question has one.

Three years after I started keeping a notebook of times I was wrong, I stopped, not because the practice wasn't working but because it was working and I'd started to believe I didn't need it anymore. I had caught so many distortions and built the habits, and the notebook felt redundant, like training wheels on a bike I could obviously ride. Six months later, I made a decision that would have been caught if I'd still been keeping the notebook: I was planning a major initiative, had consulted stakeholders, reviewed data, built a theory of change, and everything looked solid. What I hadn't done was make my predictions explicit. I hadn't written down what I expected to happen, so when early implementation went differently than anticipated, I couldn't see the gap. I kept explaining the differences as implementation challenges rather than what they actually were: flawed assumptions. If I'd been keeping the notebook, I would have had a record, the prediction next to the outcome with the separation between them visible on the page. But I'd stopped, and the groove I'd learned to catch had quietly deepened again.

I share this because the pattern repeats. I've watched it repeat in my own practice and in the leaders I work alongside. The grooves are features of human cognition. They regenerate when you stop attending to them. Staying in practice isn't a phase you complete on the way to mastery. It is the work itself. And the reduction in the time between the groove firing and the catching, the narrowing of that gap over months and years of practice, is the only honest measure of progress.

Three threats to sustained practice deserve naming.

The first is success. When you catch distortions and make better decisions, you build confidence that feels earned because it is earned. But confidence left unexamined becomes the soil where complacency grows. I've been using SEE for years and I still make brilliant blunders, less often and less severely, but the newer versions are harder to catch because they're subtler. Early on, the grooves were loud; I could hear them if I listened. Now they've learned to whisper. The patterns adapt, and what was once your blind spot becomes visible while a new blind spot forms somewhere you weren't watching.

The second is busyness. Leadership is relentless, and the practices that aren't urgent get dropped first. SEE takes time: the thirty-second version fits into the flow, but the deeper practice requires space, and when you're overwhelmed, the space feels like a luxury you can't afford. It isn't. The time invested in catching distortions prevents the much greater time spent managing the consequences of distorted decisions. But knowing this doesn't make the practice easier to maintain; you have to build it into your structures.

The third is isolation. You can do this work alone, but it's harder than it needs to be, because your own assessment of your thinking is shaped by the same grooves that distort the thinking in the first place. You need other people to see what you can't. I know a group of four district leaders who meet monthly, bringing real decisions, current and unresolved, and running SEE together. One of them told me something that stuck: "I can tell myself I've considered what I'm not seeing. But when someone else asks 'Who would disagree with this and why haven't you talked to them?', I can't dodge." The partnership works partly because there's no political entanglement; the colleague has no stake in her being right, so the feedback is clean.

Do you have people who will tell you when you're not seeing clearly? If you don't, building that may be the most important thing you can do for your practice.

The notebook has grown over the years. My original version tracked times I was wrong: the decision, what I thought would happen, what actually happened, and what the gap revealed about my thinking. Now I track predictions before major decisions with specificity: "I expect this initiative

will produce measurable changes within six months," or "I predict the board will approve this proposal with minimal concerns," or "I assume families will engage with this communication system." The specificity matters because vague predictions can't be falsified. "Things might be challenging" is always true. "I predict the implementation will take longer than we've scheduled" is testable.

I also track attribution patterns now. When something goes wrong, I write down my initial explanation and then challenge it: what if the cause is something else, what if I'm part of the cause, what would the people most affected say is driving this? Over time, my patterns have become visible: I tend to underestimate how long things take, overestimate how much teachers can absorb in a single professional development session, and attribute resistance to fear when it's sometimes about legitimate concerns I haven't addressed. Knowing my patterns helps me compensate.

The notebook is a mirror that doesn't let me forget what I've seen. Without the record, memory distorts: I remember my successes more clearly than my failures, remember predictions that came true and forget ones that didn't. The pen forces a pace the keyboard doesn't.

Every January, a superintendent I've worked with writes a letter to herself about what she learned in the previous year, not accomplishments but learnings: distortions she caught, distortions she missed, how her grooves have evolved. She also writes predictions for the coming year: where she thinks she's most likely to get caught, what decisions are coming where her grooves will be most active. Then she seals the letter and doesn't open it until the following January.

"Reading last year's letter is humbling," she told me. "I always think I've grown more than I have. I always think I've escaped grooves I'm still stuck in. The letter shows me the gap between who I think I am and who I actually am."

That gap is information about where the work still needs to happen, and it prevents the comfortable illusion that past growth means current clarity.

Here's what I want you to understand about staying in practice: it's about design rather than discipline.

Discipline fails. You cannot willpower your way to sustained practice when the demands of leadership overwhelm your attention and the grooves reassert themselves under load. Design succeeds: build the practice into your systems, schedule the reflection, create the accountability structures, join or form a practice community, keep the notebook where you'll actually use it.

The question isn't whether you will have the discipline to maintain this. The question is what structures you can create that will hold the practice even when your discipline fails.

**Calendar blocks**: scheduled time dedicated to running SEE on recent decisions. One principal blocks two hours every other Friday, reviewing what he was thinking, what he might have missed, what he would do differently. "The first few times, I didn't have much to say. But I stayed in the room for the full two hours. After a while, I started finding things."

**Decision logs**: a systematic record of significant decisions with the date, what was decided, three sentences about the reasoning, reviewed monthly. One leader noticed she was making most personnel decisions on Fridays, exactly when she was most depleted.

**Accountability partners**: someone who will ask the hard questions, who has permission to point out when you're not seeing clearly, who checks in regularly. The partnership works when there's no political entanglement and the feedback is clean.

**Trigger rituals**: "Before I finalize any hire, I write down what concerns have been raised and which ones I've dismissed." The ritual holds the practice even when memory doesn't.

**Post-mortems**: structured reviews after major decisions, asking what was predicted, what occurred, where the gap was, and what the gap reveals about how the team was thinking. Klein (1998) makes a compelling case that expertise develops from experience paired with clear feedback; without the feedback loop, years of practice can simply entrench the wrong patterns.^1^

The specific practices matter less than the principle: what will hold your practice when your attention wanders?

I want to address what you might be feeling right now. This book has asked a lot: self-examination, structural changes, ongoing practice, vulnerability with colleagues, the commitment to keep working even when the grooves adapt and the discomfort doesn't fade. If you're overwhelmed, that's a reasonable response.

You don't have to do everything at once. Start with one practice, one question you ask regularly, one structure you build, one relationship that holds you accountable. The leaders who sustain this work over years are the ones who started small and stayed consistent.

Start a decision log this week, a simple document with four columns: what you're deciding, what you expect to happen, what might go wrong, and what actually happened (filled in later). Log every significant decision for the next month, then review, comparing predictions to outcomes, paying particular attention to the concerns you raised and dismissed.

The harder version: share the log with someone else, someone who will review it with you, ask uncomfortable questions, notice patterns. If you're already planning to skip this exercise, notice that. The resistance is the groove defending itself, thriving in the absence of feedback, in the gap between prediction and

CONCLUSION

# The Students Are Waiting

A ninth grader named Dion sits in the back row of a biology class, drawing in the margins of his notebook. He is not drawing cells or mitochondria but detailed sneaker designs with shading and perspective that suggest genuine spatial intelligence. His teacher has noted this on three consecutive progress reports: "Dion is disengaged and off-task during instruction." The comment travels upward, reaching his counselor, who mentions it to the assistant principal during a hallway conversation about students who "aren't buying in." By the time the conversation reaches the leadership team's Wednesday data meeting, Dion has become a data point in an attendance and engagement spreadsheet, one of fourteen students flagged for intervention. No one in that meeting has asked Dion what he is drawing or why. No one has noticed that his test scores in geometry are higher than his scores in any other subject. No one has wondered whether the sketching might be connected to how he processes spatial information, or whether the disengagement they are tracking might be a response to instruction that has not once, in four months, asked him to use the kind of thinking he does naturally.

The decisions being made about Dion are not malicious. They are grooved. His teacher's evidence groove selects for a particular kind of engagement: verbal participation, notes in the expected format, eyes forward. Dion's engagement does not match that template, so it registers as absence rather than difference. His counselor's attribution groove locates the problem in Dion, "not buying in," rather than in the instructional environment that has given him nothing worth buying. The leadership team's prediction groove constructs a trajectory: disengaged ninth grader becomes chronically absent tenth grader becomes dropout statistic. The prediction feels like data-driven planning; it is a story they are telling about Dion's future based on a version of his present they have not examined carefully enough. And underneath the data and the categories and the intervention plans, something that no one in the room can name: the physical recognition that Dion is a certain kind of student, a sense that arrived in the body before it reached the analysis, that shaped which evidence the analysis would find, that produced the certainty that intervention is needed without ever producing the curiosity about what Dion is actually doing when he draws.

This is what the grooves cost. Not in the abstract, not in the aggregate, but in the particular experience of one student whose potential is being translated into a deficit narrative by people who genuinely believe they are helping. The translation happens in seconds, without malice, without awareness, and with total confidence that the evidence supports the concern. It does support the concern, but the evidence was selected by the same grooves that shaped the concern in the first place, and that circularity is invisible to everyone in the room.

What would change if someone in that Wednesday meeting asked the three questions?

*What am I not seeing?* Maybe someone would notice that the engagement data captures a narrow band of behavior and misses others entirely. Maybe someone would ask what Dion is doing when he is "off-task" and whether the answer complicates the story they have built.

*Who am I blaming?* Maybe someone would notice that every explanation they have generated locates the problem inside a fourteen-year-old, and that none of them have asked what the classroom is offering him or failing to offer him.

*What am I assuming?* Maybe someone would surface the prediction, "this trajectory leads to dropout," and ask what evidence supports it beyond pattern recognition. Maybe they would run a pre-mortem on their intervention plan and discover that it addresses Dion's behavior without touching the conditions that produce it.

Maybe. Or maybe not. The grooves are deep, and three questions asked in thirty seconds do not guarantee anything.

I think about the hiring story I told in the Introduction. Melissa and Terri. Two finalists for an equity coordinator position, and a committee, my committee, that selected the one who sounded like us. I ran a process I believed was fair. I asked for diverse perspectives. I summarized points of consensus. I invited dissenting views. Beneath all that procedural fairness, my grooves were building a machine that could only produce one result. My evidence groove selected for polish and fluency with theory. My attribution groove explained Terri's different interview style as a deficit rather than a difference. My prediction groove constructed a future where Melissa would succeed because she matched the template I associated with success. And underneath all of it, the sensation in the body I described in the Introduction: something in me recognized Terri as the stronger candidate, recognized it in the body during her interview, and the recognition never made it into the analysis because the analytical

track was louder and more rewarded by the culture of the room. The machine worked perfectly. It produced exactly what it was designed to produce.

I didn't tell you what happened to Terri. After we hired Melissa, after Melissa resigned eight months later, Terri had taken a position in another district. She wasn't available when we reopened the search. I reached out anyway, not to recruit her but to apologize. "I ran a process that filtered you out," I said. "I didn't see it at the time, but I see it now. I'm sorry."

Terri was more gracious than I deserved. "The system works the way the system works," she said. "You weren't trying to filter me out. You were just running your normal process."

She was right. My normal process was the problem. The grooves I had built over years of leadership were the problem. And I hadn't known it because knowing would have required me to question the very patterns I was using to evaluate everything else, including the patterns themselves. That circularity is the logic trap in its purest form.

I can't undo that decision. I can't give Terri back the position she should have had or give those teachers the leader they needed. What I can do, what any of us can do, is build the discipline of catching the pattern faster, not eliminating it because the grooves are too deep for that, too reinforced by years of habit and identity and institutional reward, but catching them sooner, more honestly, with less certainty that we have already accounted for our own distortion.

The three questions fit in your pocket. What am I not seeing? Who am I blaming? What am I assuming? You can ask them in every meeting, every conversation, every decision. They won't make you perfect. Your grooves will still fire, still construct evidence confirming what you already believe, still attribute outcomes to causes protecting your framework, still predict futures extending your assumptions rather than challenging them. But you will catch yourself more often. You will build structures that make the catching systematic. And sometimes, when it matters most, you will interrupt the groove before it determines your choice.^1^

That "sometimes" is not a small thing. For the student sitting in the back row drawing shoes, the difference between a leader who catches their grooves and one who doesn't is the difference between being seen and being sorted, between having your potential recognized and having it translated into someone else's deficit category, between an intervention that addresses what you actually need and one that addresses what the system finds convenient to target.

I haven't figured this out. I still feel certainty in meetings, still construct explanations that protect my framework, still predict futures based on what I already believe and mistake those predictions for analysis. The grooves are still there, doing their work on me every day. What has changed is that I catch them faster, not always, not reliably, but faster than I did when I sat in that conference room and selected Melissa with total

confidence. And the space separating what I sense about a situation, the felt quality that registers in the body before the analysis begins, and what my analysis says about it has narrowed enough that sometimes the felt signal interrupts the analytical process before the process has finished building its case. That narrowing is the practice. That narrowing is the work.

This is why we practice. The students sitting in our schools right now deserve leaders who are doing the work of catching themselves, leaders who have built the discipline of asking what they are not seeing, examining who they are blaming, and testing what they are assuming. Not perfect leaders. Practicing ones.

There are students right now whose futures depend on decisions their leaders haven't made yet, decisions that will be shaped by grooves those leaders may not know they have. Those students are waiting. Not for us to arrive at some finished version of ourselves. Just for us to keep practicing.

That is enough to start with.

*The Interior Architecture of Transformation continues in Book Two: PROJECTING PROOF*

^1^ Kahneman, D. (2011). *Thinking, Fast and Slow*. Farrar, Straus and Giroux.

*The Interior Architecture of Transformation*

*Joshua T. Fraser, Ed.D.*