CONTEXT IS THE JOB

Why the best HR Practitioners read the room before they reach for the playbook

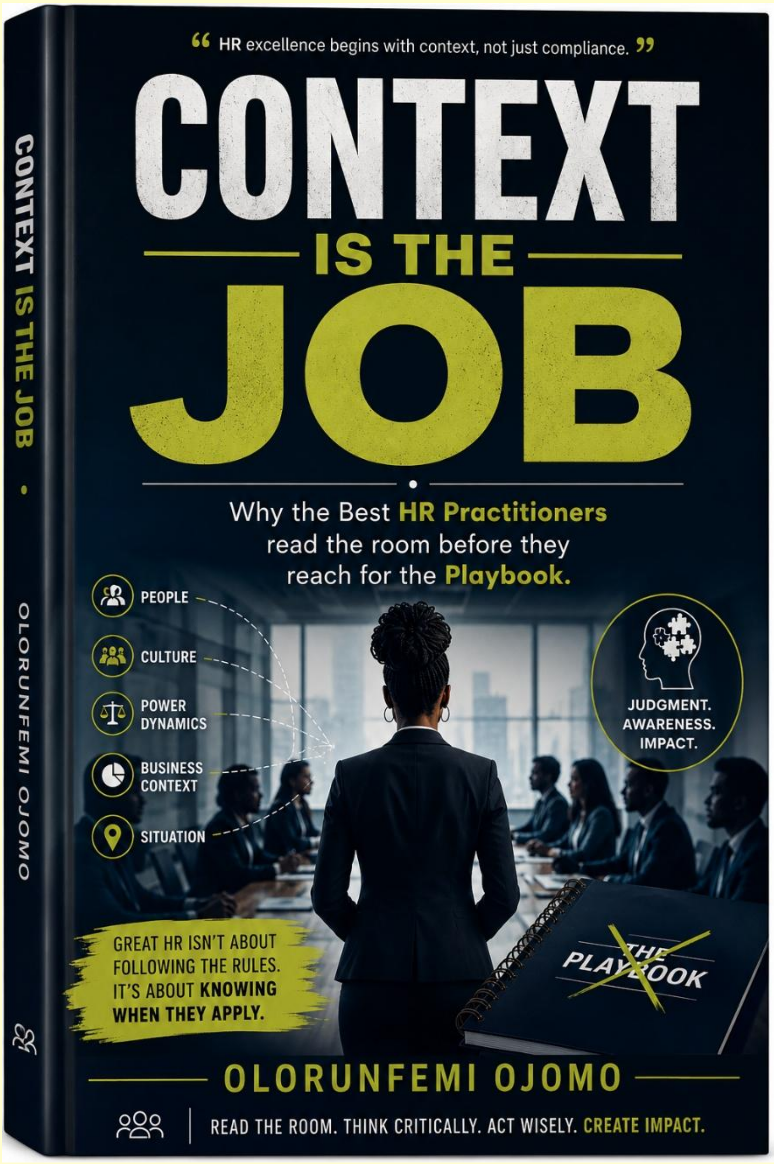

Concept by Olorunfemi Ojomo

Written using Generative AI to consolidate thought process and work within timeline. Also explored a different writing style from my conventional approach. Your feedback will be highly appreciated by sending a mail to hi@insightswsc.com.

Executive Publisher is InsightsWSC

CONTENTS

A NOTE ON THE CASES IN THIS BOOK

Every story in this book is true in the way that matters and false in the way that protects the people who lived it.

The cases that follow are composites, built from patterns observed across many organizations, industries, and practitioners over many years, not transcripts of any single identifiable company or person. Names have been invented. Industries have been changed. Timelines have been compressed, extended, or reordered. Where a detail was specific enough to point back to a real organization, it was altered or removed, even at the cost of a slightly less vivid sentence. No rating scale, severance formula, or org chart in this book belongs to a real company, even where it is described with enough precision to feel like it does.

This isn't a legal hedge dressed up as a literary one, though it is certainly that too. It reflects something closer to the argument of the book itself. The value in these stories was never in their specificity as gossip; who, exactly, at which company, said what, to whom. The value is in the pattern: the shape of a technically sound tool landing badly on a room it wasn't built for. A composite can carry that pattern faithfully, often more faithfully than a single real case could, because a single real case always carries some idiosyncratic noise alongside the signal, and a careful reader deserves the signal without having to guess which parts of the story to trust.

A small number of cases in this book: one per Part, generally the most detailed and most narratively developed are more heavily reworked versions of situations closer to real events: industry changed, scale altered, sequence rearranged, until the people who actually lived through them would not recognize themselves in the retelling, even as the underlying pattern remains intact. These are marked by nothing more than the level of detail in the telling; there is no asterisk system distinguishing them from the lighter composites, because the distinction matters less to the reader than the honesty of the method behind all of it, which is stated here, once, plainly, and applies to every case in the book without exception.

If any story in these pages feels uncomfortably familiar, that is very likely because the pattern is more common than any of us would like to admit, not because the story is about you, or your company, or anyone in particular. That familiarity is, if anything, the point.

AUTHOR'S NOTE

This book came to be after I overheard a senior practitioner trying to pass judgement on the interventions of a young developing practitioner through name dropping of models which is deemed fit as best practice in the subject of discuss without recourse to the context in which the young professional played. This situation has always played out and continually pile up conviction for me to write about it.

There is a list of things we wished someone had told us earlier; it is not about how to run a compensation study or facilitate a values workshop, since those things were covered, thoroughly, by the training and certifications that got us into this field in the first place. It started as a list of the moments, where doing everything by the book still went badly, and where the things that would have actually helped had no name within the league of models and best practices.

Every practitioner reading this have their own version of that list. Some of the situations in this book will be familiar because something close to them happened at a company you worked for. Others will be familiar for a stranger reason: because the shape of the failure matches something you lived through even though none of the surface details do. That's the pattern this book is trying to name, and if it succeeds at anything, it will be at giving practitioners a shared vocabulary for a kind of judgment most of us have been quietly exercising, unevenly and without much support, for as long as we've been doing this work.

This book does not claim to replace the technical foundations of HR practice, and it isn't meant to. Compensation design, employment law, performance methodology, structured hiring; these remain genuinely necessary, and nothing in the chapters ahead should be read as an argument against rigor in any of them. The argument is narrower and, I'd suggest, more useful than a wholesale critique of the field's technical core: that rigor in the tool has to be matched by rigor in reading whether, when, and how to use it, and that second kind of rigor has never quite gotten the same attention as the first.

If this book has a single hope for the reader, it's a modest one. Not that you'll memorize five letters and a diagnostic checklist, though I hope PACES earns a place in your actual practice. It's that the next time something technically correct goes quietly wrong, you'll have a faster route to the real question than most of us had for most of our careers; not "what did I do wrong with the tool," but "what did I not yet understand about the room."

That question, asked early enough and honestly enough, is most of the job. The rest, as the closing chapter puts it, is vocabulary.

PROLOGUE: THE ROOM SHE DIDN'T READ

The exit interview lasted eleven minutes, and the reason she was leaving took up about four seconds of it.

“It's not the ratings”, the lady said, glancing at the door like she was already somewhere else. “It's just that we've heard this before.”

Maya wrote it down. She wrote most things down; it was one of the reasons she had been hired. But she didn't understand the sentence yet, not really, and she moved on to the next question on the form, because that's what exit interviews are for: closing files, not opening them.

She would think about that sentence for the next four months.

Maya Okonkwo had done this twice before, at two other companies, and both times it had worked. She knew what she was doing. Calibrated ratings, a five-point scale collapsed down to three to stop the grade inflation. Structured feedback conversations, twice a year, with a rubric so managers couldn't hide behind vague adjectives. A promotion process with actual criteria attached to it; not a black box, not a tap on the shoulder from someone's favourite GM. It was, by any reasonable industry standard, a good system. Modern. Fair. The kind of thing that shows up in the same conference talks year after year because it keeps working.

She rolled it out at her new company with the confidence of someone who had watched this exact movie twice and knew how it ended.

The first quarter went fine. Managers grumbled about the training time, which managers always do. A few people asked pointed questions about how ratings would map to compensation, and she'd answered them cleanly, because she'd answered them before. There was, if anything, a kind of relief in the room; finally, something that looked rigorous, something that looked like the company was growing up.

It was around month five that she noticed the pattern in the attrition data. Not a wave. A rise. A rise made entirely of people she didn't want to lose; the ones with the highest ratings, the ones managers fought to keep on their teams. Engagement scores hadn't moved. Compensation was competitive; she'd checked twice. Managers described the new process as clear, even good. Nobody could tell her why the best people kept finding reasons to update their LinkedIn.

She started sitting in on more exit interviews herself.

That's where she heard it the second time; a different employee, a different phrasing, the same shape. "I don't think the system's the problem," he said. "I think everyone's just waiting to see how long this one lasts."

This one.

She asked him what he meant, and he shrugged, the particular shrug of someone who assumed she already knew and was being polite by not making her say it out loud. So she asked someone else. A team lead, five years tenured, the kind of person who had outlasted three transformation and never once shown up in an exit interview. He didn't shrug. He told her.

This was the third performance management system the company had tried in four years.

The first one had been introduced with real fanfare; an offsite, a keynote from the CEO, a promise that this time, ratings would finally connect to real money and real growth paths. It lasted five months before a reorg swallowed it whole, and every rating that quarter simply evaporated, unacknowledged, undiscussed. The second one arrived eighteen months later, quieter, more cautious, clearly built by someone trying to learn from the first failure. It made it eight months. A new VP of Sales didn't like being told what to do by a rubric, pushed back hard enough in a leadership meeting, and the whole thing was quietly shelved. Nobody sent an email announcing its death. It just stopped being mentioned.

And now there was a third. Rigorous. Fair. Modern. Exactly the kind of system that should have worked, because on its own terms, it was good.

Except nobody in that room was evaluating it on its own terms. They were evaluating it against a four-year pattern that said: systems here don't survive contact with the organisation. The best people; the ones with the most options, the ones least inclined to spend a fourth year waiting to find out if this cycle would matter had simply done the math and left before finding out.

Maya hadn't done anything wrong. Every individual decision she'd made was defensible, and in a different company, on a different history, this exact framework, delivered this exact way, would likely have worked exactly as it had the two times before. That was the part that unsettled her most, once she finally understood it. It wasn't a competence problem. She hadn't failed to execute the framework.

She'd failed to read the room the framework was landing in.

This is not a book about performance management systems, and it is not really a book about Maya, who is a composite of several people this has happened to, in several rooms much like this one. It's a book about the four seconds in that first exit interview; the ones that got written down but not understood, because nothing in her training had taught her to listen for them.

Here is the sentence that took Maya four months to hear:

It's not the ratings. It's just that we've heard this before.

Every practitioner who has been in this field longer than a year has a version of that sentence somewhere in their own history said to them directly, or said behind their back, or never said at all because by the time someone might have said it, the damage was already done and nobody saw the point.

The premise of this book is that those four seconds are the real job. Everything else: the frameworks, the rubrics, the rollout plans is just the vocabulary. Context is the grammar. And most of us were only ever taught the vocabulary.

PART I - THE BLIND SPOT

CHAPTER 1: THE MYTH OF BEST PRACTICE

The consultant's slide deck was very good. That was, in some ways, the whole problem.

Halvorsen Steel had been making rebar and structural components in the same Ajaokuta valley for sixty-one years, run by three generations of the same family, employing just under four hundred people, most of whom had a cousin or a father who'd worked there before them. The plant manager still walked the floor every morning and knew which machines to listen to for trouble before the sensors caught it. When the third-generation CEO decided the company needed to "modernize how we set goals," he hired a firm that had done exactly this kind of engagement for a dozen tech companies, and the firm did what firms do: it brought the playbook that had worked at the last dozen companies.

OKRs. Objectives and Key Results, the goal-setting system built at Intel and popularized, spectacularly, at Google; quarterly cycles, ambitious "moonshot" targets deliberately set high enough that hitting seventy percent was considered a win, cascading visibly from company-wide objectives down through every team, reviewed and re-graded in public at the end of each quarter. It is, by the standards of the world it came from, an excellent system. It rewards ambition. It creates alignment across teams that don't naturally talk to each other. It has a real, well-documented track record.

None of that mattered once it landed on the floor at Halvorsen.

The plant's key results were things like reducing scrap rate by an ambitious, deliberately uncomfortable twenty percent; a number chosen precisely because the OKR philosophy says stretch goals should feel slightly out of reach. To a product team at a software company, a quarter where you hit seventy percent of an aggressive target is a good quarter, worth celebrating. To a floor supervisor who had spent eighteen years being evaluated on whether a batch met spec, "seventy percent of a stretch goal" read as one thing: a fail, dressed up in language designed to obscure it. Nobody had explained the philosophy, because the consultants assumed, correctly by their own experience, that everyone already understood how OKRs worked. Nobody at Halvorsen did.

The quarterly review meetings, imported wholesale from the tech playbook, called for teams to grade themselves publicly on a red-yellow-green scale and discuss what had gone wrong. In a Bay Area engineering org, this ritual is one of psychological safety; everyone's OKRs are ambitious, everyone’s partly red, nobody is singled out. At Halvorsen, where evaluation had always been private, quiet, and handled by a supervisor who knew you, the public grading felt like something closer to a tribunal. Men who had worked the same line for two decades sat through meetings watching their names attached to red scores in front of people they'd trained. Two long-tenured supervisors quietly asked to be moved off the shop floor within the quarter. A third, one of the plant's most respected hands, retired eight months earlier than he'd told anyone he planned to.

The CEO, to his credit, noticed the damage before it became catastrophic and killed the program after two quarters. But the trust cost lingered for years; every subsequent initiative from corporate, however sensible, was met with a wariness that hadn't existed before. "Here comes another one of those," became something people said under their breath in the break room, and everyone knew what "those" meant.

Nothing about OKRs is inherently bad. The system is precisely as good as it was ever claimed to be - for the context it was built inside. Google built OKRs for Google: a company selecting for high-ambition, high-autonomy engineers, evaluated already on individual output, operating in a culture that treated public failure as informational rather than personal. Halvorsen Steel was none of those things, and no one asked whether it was before installing the same operating system.

This is the quiet fraud sitting inside the phrase "best practice." Best is a comparative word. It requires an answer to *best compared to what, for whom, under which conditions* and the phrase, as it's used in almost every HR conference talk and vendor pitch deck, simply drops that half of the sentence and hopes nobody notices. A practice cannot be best in the abstract, the way a screwdriver cannot be the best tool in the abstract. It is only ever the best tool for a particular screw, in a particular material, at a particular angle. Call OKRs the best goal-setting practice and you've made a category error before you've even opened the deck; you've described a *tool* using the grammar of a *universal law*.

The phrase has a real history, and it's worth knowing, because the history explains why it survives despite being, on close inspection, almost meaningless. "Best practice" entered mainstream management vocabulary through the benchmarking movement of the 1980s, popularized heavily by firms like McKinsey and later formalized by bodies like APQC, built on a genuinely useful idea: study what high-performing organizations do differently, and learn from it. The trouble is what happened on the way from that idea to a PowerPoint slide. Benchmarking studies compare practices across companies with wildly different sizes, industries, cultures, ownership structures, and histories, and then compress the findings into a tidy, portable recommendation because a tidy, portable recommendation is what sells at a conference and fits on a single slide with a checkmark icon next to it. The context that made the practice work in the original company gets quietly edited out somewhere between the research and the readout, the way a photograph loses its background when you crop it down to a headshot.

By the time a "best practice" reaches an HR practitioner at a family-owned steel manufacturer in Ajaokuta, it has been stripped of everything that would tell her whether it applies to her. What's left is a claim of universality that was never true even in the place it came from.

This is not an argument against learning from other organizations, and it's not an argument that OKRs, or nine-box grids, or any other widely used HR tool is secretly bad. It's an argument that the sentence "this is a best practice" should trigger suspicion rather than comfort in anyone who hears it, because the sentence is doing exactly the opposite of what good practice actually requires. Good practice requires asking what conditions this worked under, and whether those conditions exist here. "Best practice," as a phrase, is specifically engineered, not maliciously, just structurally, by the incentives of the industry that produces it to let you skip that question.

The next chapter asks a harder one: if it isn't a competence problem and it usually isn't, what is it?

CHAPTER 2: COMPETENCE IS NOT THE PROBLEM

It is worth being precise about what did not go wrong at Halvorsen Steel, because the precision matters more than it might seem to.

The consultants who designed the OKR rollout knew OKRs. They could have taught a masterclass on cascading objectives, on the difference between a good key result and a vanity metric, on how to run a quarterly grading session. Their technical command of the tool was not in question, and it would be a mistake to read the last chapter as a story about people who didn't know what they were doing. They knew exactly what they were doing. They simply didn't know where they were doing it.

That distinction between knowing the tool and knowing where the tool is landing is the one most HR training never draws, because most HR training is built entirely around the first half. Certification programs teach you to run a compensation study, structure an interview loop, calculate a nine-box distribution, facilitate a values workshop. These are real skills, and they take real years to build. Call this *technical competence*: the ability to execute a given HR practice correctly, according to its own internal logic, the way a textbook or a certifying body would grade it.

There is a second kind of competence that almost nothing in the field explicitly trains for, and it is the one that actually determined what happened at Halvorsen. Call it *contextual competence*: the ability to read an organization accurately enough to know whether, when, and how a given practice should be used at all, and, just as often, the judgment to know when a technically correct practice should be delayed, modified beyond recognition, or abandoned entirely, because the room it's about to enter isn't ready to receive it in its textbook form.

These two competencies are not the same skill wearing different clothes. They are genuinely different capacities, and a practitioner can have a great deal of one and almost none of the other. This is uncomfortable to sit with, because HR as a field has spent decades building its professional identity around the first kind: certifications, credentials, frameworks, a whole apparatus for proving you know the tools while treating the second kind as something closer to instinct or personality, a soft "people sense" that some practitioners just happen to have, rather than a discipline that can be named, studied, and deliberately built.

This is the sentence worth sitting with for the rest of the book: most HR failures are context failures, not competence failures. Not incompetence: context. The rollout that alienates a workforce, the DEI program that triggers backlash instead of buy-in, the performance system that tanks retention, the culture initiative that becomes an office joke within a quarter, in the postmortem, these are almost never traced back to someone not knowing how the tool works. They're traced back to someone applying a correctly-executed tool into a room that wasn't the room it was built for, at a moment that wasn't the moment it needed, in a way that ignored what the room already believed.

This matters enormously for how organizations respond when an HR initiative fails, because the two failure types call for opposite remedies, and most organizations reach for the wrong one by default. A competence failure calls for more training, better certification, a more rigorous process, a more experienced hire. A context failure calls for something almost nobody puts in a job description: slower diagnosis before action, deliberate listening before rollout, and most uncomfortably, the humility to sometimes not use the tool you know is technically correct, because the room isn't ready for it yet.

When organizations misdiagnose a context failure as a competence failure, the fix makes things worse. The company brings in a more credentialed consultant, who runs the exact same technically excellent playbook with even more polish and confidence, and the room reads the extra polish as extra tone-deafness rather than extra rigor. This is a very common pattern, and it's worth naming plainly: many HR functions have spent years cycling through increasingly qualified people to solve a problem that has nothing to do with qualification. The revolving door of Heads of People at a company with a genuinely toxic history isn't evidence that HR keeps hiring the wrong person. It's often evidence that the organization keeps solving a context problem by throwing more competence at it.

There is a reason contextual competence gets so little formal attention, and it isn't that the field has overlooked it by accident. Technical competence is legible, it can be tested, certified, put on a resume, demonstrated in an interview with a case study. Contextual competence is almost entirely illegible by comparison. It shows up as a hesitation before recommending something. A decision to run three more conversations before drafting the rollout plan. A willingness to tell a CEO that the technically best option is the wrong one for this year. None of that fits neatly into a certification exam, so the field has, understandably if not forgivably, built its entire training apparatus around the half of the job that can be measured, and left the other half to be picked up informally, unevenly, usually the hard way, usually after a Halvorsen-sized mistake of one's own.

The chapters that follow aim to make contextual competence considerably more legible than that; to give it a shape, a vocabulary, and eventually a framework practitioners can actually use before the rollout, not just recognize in hindsight after the exit interviews start. But first, it's worth understanding why the field keeps reaching for the wrong half of the job in the first place. That's not an accident either. It's built into the structure of the industry itself.

CHAPTER 3: WHY HR KEEPS IMPORTING SOLUTIONS

Ask any experienced practitioner where their favourite framework came from, and the honest answer is rarely "I built this after diagnosing our specific situation." More often it's a conference session, a webinar, a case study a vendor presented at a regional SHRM or CIPM chapter meeting, or a slide a consultant left behind after an engagement at a different company entirely. This isn't a personal failing. It's the predictable output of an industry whose major institutions are, almost without exception, structured to reward having a framework over fitting one.

Start with the conference circuit, because it's the most visible. HR conferences; the large national ones and the smaller regional chapters alike are built around a specific kind of content: the case study talk, in which a practitioner from a well-known company walks an audience through "how we did X," usually X being some initiative that produced a good, quantifiable result. These talks are genuinely useful as far as they go, and the practitioners giving them are rarely being dishonest. But the format itself has a structural flaw: a forty-five-minute conference session has no room for the six months of context-gathering that preceded the initiative, the two failed attempts before this one, the specific political conditions inside that particular company that made the timing work, or the version of this initiative that quietly failed at a sister company down the street. What survives the editing process, by necessity, is the transferable-sounding part; the framework, the five-step model, the acronym because that's the part that fits on a slide and photographs well for LinkedIn afterward. The audience leaves with the tool and none of the conditions that made the tool work, and there is no format at a conference that naturally corrects for this, because the format itself is optimized for portability, not accuracy.

Certification bodies compound the problem in a different way. This isn't a criticism of certification as a concept; a shared professional vocabulary and a baseline of technical knowledge genuinely matter, and the practitioners who go through this training are almost always better for it. But certification exams, by their nature as exams, need answers that are gradable, and "it depends on the organization" is not a gradable answer. So, the material that gets tested tends to be the material that has a right answer: the formula, the compliant process, the textbook definition of a nine-box grid. Contextual judgment; the skill of knowing when the textbook answer is the wrong one is nearly impossible to write a multiple-choice question around, so it quietly disappears from the curriculum, not through anyone's decision to exclude it, but because the exam format was never built to hold it in the first place.

Then there's the vendor sales cycle, which has the least excuse of the three, because unlike conferences and certifications, it operates on a genuinely misaligned incentive rather than an accidental one. A vendor selling a performance management platform, an engagement survey tool, or an AI-driven talent system is not incentivized to tell a prospective client that their organization isn't ready for it, or that the tool is wrong for a company at their stage, or that the real problem is something the software can't touch. The sales cycle is built around demonstrating fit, not diagnosing its absence; the demo is designed to make the tool look universally applicable, because a demo that says "this depends heavily on your context, let's find out if you qualify" doesn't close deals. HR practitioners, who are often under real pressure to show measurable progress on some initiative, become the internal champions for tools that were pitched to them as broadly, confidently applicable, and the vendor has no structural reason to correct that impression before the contract is signed.

The last piece is the least discussed and possibly the most consequential: the sheer speed of CHRO and Head of People turnover at many organizations, particularly ones already struggling. A new HR leader arriving at a company under pressure to show results quickly faces an obvious incentive problem. Genuine contextual diagnosis; actually understanding an organization's history, its power structure, its scar tissue from past initiatives takes months of quiet listening that produces very little visible output in the short term. A recognizable, credentialed framework, on the other hand, can be announced in a town hall within the first ninety days, generates an immediate sense of momentum, and photographs well in a board deck. A leader whose tenure might reasonably be expected to last eighteen months to two years, a pattern common enough in struggling organizations to be its own well-documented phenomenon, is rationally incentivized to import a recognizable solution fast rather than spend a quarter of their limited runway diagnosing a context that a successor might not even inherit the credit for understanding. This isn't a story about poor character. It's a story about an incentive structure that punishes exactly the kind of patience contextual competence actually requires.

Put these four forces together, a conference format that can't carry context, a certification system built to test what's gradable, a vendor economy incentivized to claim universal fit, and an executive tenure cycle too short to reward patient diagnosis and the result isn't a mystery. It's the entirely predictable output of an industry where every major institution, for reasons that mostly make sense from the inside, rewards a practitioner for arriving with a framework, and does almost nothing to reward the slower, less visible work of finding out whether it fits.

None of this is presented as an indictment of any individual practitioner reading this book. Almost everyone in this field has, at some point, reached for the framework from the conference talk instead of doing the harder diagnostic work first; because the entire apparatus around this profession was built to make that the path of least resistance. The point of naming it this plainly isn't to assign blame. It's to make the pattern visible enough that it stops feeling like personal failure and starts feeling like something that can actually be corrected, which is where the rest of this book picks up.

PART II - THE GRAMMAR OF CONTEXT

CHAPTER 4: A NAME FOR WHAT YOU ALREADY DO

Every experienced HR practitioner already does some version of what this section is about to formalize. Before a rollout, most competent people ask a handful of questions almost by instinct; is this the right time, will the CEO actually back this if it gets uncomfortable, has something like this been tried before. The problem isn't that this instinct doesn't exist. It's that it stays an instinct: informal, inconsistent, easy to skip under deadline pressure, and nearly impossible to teach to someone else, because "just read the room" is not instruction, it's an outcome dressed up as advice.

This section gives that instinct a shape. Five dimensions, run as a diagnostic before any tool gets chosen, in roughly the order they tend to reveal themselves in a new engagement: Power, Ambient culture, Conditions, Ego, Stage. Together they spell PACES, which is more than a mnemonic convenience; pace is precisely the right word for what this framework is trying to teach. An organization has a pace of its own: how fast it can absorb change, how much discomfort it can metabolize before it protects itself, how quickly trust rebuilds after it's spent. Reading PACES correctly means reading that pace before setting yours, rather than importing a rollout timeline built for someone else's organization and its own rhythm.

A brief note on the arithmetic, because a careful reader will notice the original list this framework is built from had six items, not five. Precedent; what happened last time something like this was tried is real, and it matters enormously, as the Halvorsen and Maya stories both showed. But precedent is not a separate force acting on an organization the way power or external pressure are. Precedent is how an organization's ambient culture remembers. It's the specific, retrievable evidence for the general climate; the receipts, not the weather. Folding it into Ambient culture rather than treating it as its own dimension isn't a simplification for its own sake; it reflects something true about how organizations actually work. You rarely encounter "precedent" floating free of culture. You encounter it as the reason a culture feels the way it does. Ask a five-year employee why the room is skeptical of a new initiative, and precedent is usually the answer they give when they're asked to explain the ambient culture question. They're the same conversation.

Each of the five chapters that follow works the same way: a definition of the dimension, why it gets missed so often, a short diagnostic; the specific questions worth asking before committing to any HR initiative and a brief illustration of the dimension in action. These aren't yet the full case studies; those are coming in Part III, where every major HR domain gets run through the same tool in two different contexts to show how differently it lands. Here, the goal is narrower and more foundational: to make each of the five lenses individually legible, sharp enough to actually use, before Part III shows what happens when all five are read together.

A caution worth stating plainly before diving in. PACES is not a checklist to complete once and file away, and it is not a scoring system that outputs a single number telling you whether to proceed. Two organizations can score identically on all five dimensions and still require completely different approaches, because the dimensions interact with each other in ways a simple checklist can't capture; a founder-led company under intense external pressure behaves very differently than a founder-led company with no pressure at all, even though "founder-led" is the same answer on the Power dimension in both cases. Treat PACES the way a pilot treats a pre-flight checklist: not a substitute for judgment, but a structure that makes sure judgment doesn't skip something it would otherwise miss under pressure or fatigue. The checklist doesn't fly the plane. It makes sure the person flying the plane has actually looked at the instruments before take-off.

One more thing worth saying before the individual chapters begin. Practitioners who take this framework seriously will sometimes find that the honest diagnosis leads to an uncomfortable recommendation: not yet, not like this, or not at all. A textbook-correct initiative, reading badly against the organization's actual power structure, culture, external pressure, leadership psychology, or stage, sometimes needs to be delayed, restructured beyond recognition, or shelved entirely; even when nothing about the initiative itself is wrong. This is the least comfortable implication of everything in this book, because it asks practitioners to sometimes advocate for doing less, or doing it more slowly, in a professional culture that often rewards visible action over invisible judgment. It is also, if the argument in Part I holds, the actual job. The five chapters ahead are meant to make that job practicable, not just persuasive.

CHAPTER 5: P FOR POWER

Every organization has an org chart, and almost none of them tell you who actually has to say yes!

The org chart shows formal authority: titles, reporting lines, the boxes and dotted lines that appear in the employee handbook. The Power dimension of PACES is asking a different, more useful question: when an initiative needs real backing to survive contact with resistance, whose actual agreement determines whether it lives or dies? In a great many organizations, the answer bears only a loose relationship to the org chart and mistaking the chart for the real power structure is one of the most common and most avoidable context failures a practitioner can make.

Consider four common power structures, because the differences between them aren't subtle once named. A founder-led company runs on the founder's continued conviction; an initiative can have every VP's sign-off and still die the moment the founder loses interest or gets a better idea in the shower, because the founder's authority was never really delegated in the first place, no matter what the reporting lines suggest. A professionally managed company, run by executives who came up through careers rather than ownership, tends to distribute real authority more genuinely across a leadership team, which means initiatives need broader coalition support but are also less vulnerable to a single person's mood. A private-equity-owned company answers to a board with its own timeline, typically compressed around an exit horizon, which means initiatives get evaluated less on cultural fit and more on whether they move a metric that matters to that timeline. A family business carries authority along bloodlines that may or may not track with formal titles at all; the person with the actual final word might be a retired founder who still has lunch with the current CEO every Tuesday, holding no title whatsoever on the current chart.

None of these structures is better or worse than the others in the abstract. Each simply requires a completely different theory of how to get something to stick. An initiative pitched and sold the way you'd sell it in a professionally managed company; build the coalition, get every functional leader's buy-in, present a business case to a committee will often stall indefinitely in a founder-led company, where that same coalition-building reads as bureaucratic hedging to a founder who wants a direct, personal yes or no. The reverse mistake is just as common and just as costly: pitching a private equity board on cultural transformation language when what they actually need is a clear line to the metric their own investors are asking about.

The diagnostic questions worth asking before any initiative, regardless of what the tool itself is, are these: If this initiative meets serious resistance six weeks in, whose continued support determines whether it survives? Is that person's authority formal, informal, or some blend the org chart doesn't capture? What has this person historically needed to see before backing something publicly: data, a peer's endorsement, a personal relationship with the person proposing it, visible enthusiasm from their own inner circle? And critically: has this person actually said yes, or has someone one level below them said yes on their behalf, assuming they'd be fine with it?

That last question deserves its own moment, because it's where a surprising number of failures originate. A well-meaning HRBP gets sign-off from a VP, builds an entire rollout plan around that sign-off, and only discovers at launch that the VP never actually cleared it with the person who really controls the outcome: a founder, a board member, a family patriarch who still shows up to strategy meetings uninvited. The initiative doesn't fail because it was wrong. It fails because the practitioner mistook a proxy's approval for the real thing, in an organization where the real thing was never actually up for delegation.

Reading power correctly doesn't mean becoming cynical about organizational structure, and it doesn't mean abandoning the standard playbook of stakeholder mapping and coalition-building; those tools remain genuinely useful. It means running them against the actual structure of authority rather than the one printed on the handbook's org chart, and being honest with yourself, early, about which one you're looking at.

CHAPTER 6: A FOR AMBIENT CULTURE

Every organization is haunted, in the mildest and most literal sense of that word by the accumulated residue of everything that has previously happened to the people in it. Ambient culture is the name for that residue: the unspoken rules, inherited memory, and this is where precedent lives; the specific, retrievable evidence of what happened the last few times something like this was attempted.

The word ambient is chosen deliberately over something more dramatic like trauma or history, because the effect is usually closer to weather than to wound. Nobody at Maya's company was actively grieving the two failed performance systems that preceded hers by the time she arrived. Nobody brought it up in her onboarding. It simply sat in the atmosphere of every conversation about the new system, shaping how people heard her announcements, without anyone needing to consciously reference it. That's what makes ambient culture so easy to miss: it rarely announces itself as an active grievance. It shows up as a slightly-too-quiet room, a joke that doesn't quite land, a shrug instead of a direct answer signal easy to read as something else entirely or not read at all.

Precedent is the concrete backbone of ambient culture, and it's worth treating as its own diagnostic question precisely because it's the part that can actually be investigated rather than merely sensed. Before rolling out anything, a new performance system, a comp change, a return-to-office policy, a reorg; a practitioner can and should simply ask: has something like this been tried here before, and what happened? This sounds almost too obvious to state, and yet it is skipped constantly, for a specific and understandable reason: the people best positioned to answer it honestly are rarely the people in the room when a new initiative is being planned. Executives sponsoring a new initiative often weren't around for the last failed attempt or have genuine incentive not to dwell on it. The people who remember, the five-year employee, the team lead who trained half the department, the admin who's outlasted four reorgs are usually several levels removed from the planning conversation, and nobody thinks to ask them until after something has already gone wrong.

This is one of the most practical, immediately usable pieces of advice in this entire framework: before launching anything, find the organization's institutional memory; not the leadership team, the people who have simply been there the longest, regardless of title and ask them a version of the precedent question directly. What has the company tried before that looked like this? How did it go? What do people around here expect to happen when something like this gets announced? The answers to these questions are frequently more predictive of an initiative's success than anything in the initiative's own design.

Beyond precedent specifically, ambient culture also includes the broader unspoken rules an organization runs on; norms nobody has written down because everyone who's been there long enough absorbed them without being told. Whether disagreement with a leader happens in the room or only in the parking lot afterward. Whether asking for help reads as appropriate or as weakness. Whether the stated value of "transparency" on the careers page actually describes how information moves or describes an aspiration the leadership team is still working toward. These norms are frequently invisible to leadership themselves, who experience the culture from a position of relative safety that insulates them from its sharper edges, and are often only visible from several rungs down.

A practitioner reading ambient culture well develops something close to an ethnographic habit: treating the organization, at least during the diagnostic phase, less like a client to be served efficiently and more like a place to be understood before being acted upon. This runs directly against the pressure most HR practitioners are under to move fast and show results, which is precisely why it gets skipped so often, and precisely why the organizations that skip it keep landing in the same pattern Halvorsen and Maya's company both fell into: a technically sound initiative, executed by a capable practitioner, defeated by a history nobody thought to ask about.

CHAPTER 7: C FOR CONDITIONS

Conditions are the forces squeezing an organization from outside its own walls; the market it competes in, the investors it answers to, the regulatory environment it operates under, the presence or absence of a union, the broader economic climate its industry is riding. Unlike power and ambient culture, which are largely internal to the organization, conditions are mostly beyond any single leader's control, which is exactly what makes them easy to underweight in HR planning: they don't show up in an org chart or an employee survey, but they shape what an organization can actually absorb at any given moment just as powerfully as anything internal does.

The clearest illustration of conditions at work is timing. The identical initiative: say, a generous but restructured benefits package that trades a small amount of guaranteed comp for more flexible perks can land as a thoughtful, well-received improvement in a company riding a strong growth year, and as a callous, badly-timed provocation at the exact same company eighteen months later, mid-layoff, under investor pressure to cut costs, with nothing about the initiative's actual design having changed. The tool didn't get worse. The conditions around it did, and a practitioner who designs entirely around the tool's internal logic, without checking the external weather, will be blindsided by a reaction that has almost nothing to do with the initiative itself.

Investor timelines deserve particular attention, because they're frequently invisible to HR practitioners who don't sit close to the board, and yet they shape leadership's actual appetite for risk more than almost anything else. A company eighteen months from a planned exit behaves very differently than the same company five years out; initiatives that improve long-term retention at real short-term cost get a much harder hearing when leadership's own compensation is tied to a near-term liquidity event. This isn't cynicism about executive motives; it's simply an accurate read of the incentive structure a practitioner is operating inside, and initiatives pitched without accounting for it tend to get quietly deprioritized for reasons that never make it into the meeting notes.

Regulatory heat and union presence work similarly, shifting what's actually possible regardless of what's technically desirable. An organization under active regulatory scrutiny, or mid-negotiation with a union, has constraints on what HR can say and do that have nothing to do with best practice and everything to do with legal exposure and negotiating leverage; a well-intentioned pay transparency initiative, textbook-correct in isolation, can complicate an active union negotiation in ways that have nothing to do with whether transparency is a good idea in principle.

The broader economic and industry climate rounds out this dimension, and it's the one most likely to change without anyone inside the organization doing anything at all. An industry-wide downturn, a sudden shift in how investors are pricing growth versus profitability, a competitor's high-profile layoff; none of these originate inside the company, and yet all of them shift what the workforce is primed to expect, and what leadership is willing to risk, sometimes within a matter of weeks.

The diagnostic questions here are less about the organization itself and more about the environment pressing on it: What is this organization's investor or ownership timeline, and how does that shape leadership's tolerance for near-term discomfort? Is there active regulatory or legal exposure that constrains what can be said or promised? Is there a union presence, active organizing effort, or recent labour action that changes the stakes of this particular initiative? What's happening in this company's specific industry right now? Is this a moment of confidence or a moment of anxiety, and does the organization's public messaging about that moment match what leadership actually believes privately?

Conditions are the dimension most likely to change the fastest and the one most practitioner check least often, largely because it requires looking outward, at markets and boardrooms, rather than inward at the organization's own people and history. A diagnosis that was accurate three months ago can be badly out of date by the time an initiative actually launches, which is a strong argument for treating the PACES read not as a one-time exercise but as something worth revisiting close to any major rollout, not just at its inception.

CHAPTER 8: E FOR EGO

This is the dimension practitioners are least likely to say out loud, and the one that quietly determines more outcomes than almost any other. Ego, in the sense meant here, isn't a moral judgment about vanity or arrogance. It's a neutral description of something every leadership team has: an actual risk tolerance, an actual relationship to being publicly wrong, an actual self-image; and a gap, often a very wide one, between that reality and the culture the organization claims to have on its careers page and in its values deck.

Every company describes itself, in its public-facing materials, in fairly similar terms: we value transparency, we're data-driven, we welcome dissent, we're a learning organization that treats failure as a growth opportunity. Some of this is genuinely true. A great deal of it is aspirational at best and directly contradicted by how leadership actually behaves under real pressure at worst. The gap between the stated culture and the leadership team's actual psychology is precisely what the Ego dimension is trying to surface, because initiatives designed around the stated culture, rather than the real one, tend to fail in a very specific and very avoidable way.

Consider an engagement survey with fully anonymous, unfiltered open-text comments; a practice that fits comfortably inside a "transparency and psychological safety" stated culture. In an organization where the executive team's actual relationship to criticism matches that stated value, this works exactly as designed: leadership reads hard feedback, absorbs it, and the survey becomes a genuine trust-building exercise. In an organization where the stated value is aspirational and the executive team's real relationship to public criticism is considerably more defensive, the same fully anonymous survey produces a very different, very predictable second-order effect: leadership sees sharp comments, feels attacked rather than informed, and the HR function that ran the survey inherits the blame for "creating negativity," regardless of the fact that the negativity was already there and the survey simply made it visible. The tool didn't create the problem. It exposed a gap between stated and actual leadership psychology that nobody had diagnosed before launch.

Reading this dimension accurately requires paying more attention to how leaders behave under pressure than to how they describe themselves, because the two are frequently disconnected without any dishonesty involved, most executives genuinely believe the version of themselves in the values deck and are describing an aspiration sincerely rather than a reality cynically. A useful diagnostic isn't to ask leadership directly how they'll react to hard feedback or public failure, since the answer will almost always describe the aspiration rather than the pattern. It's to look at how they've actually reacted the last several times something uncomfortable surfaced: a bad survey result, a public misstep, a piece of negative press, a direct challenge from a subordinate in a meeting. Did they get curious, or did they get defensive and start looking for who was responsible? Did the messenger get protected, or did the messenger quietly stop getting invited to certain meetings? Patterns like these are far more predictive than anything a leader will say about themselves in an interview, because they describe what actually happened rather than what someone believes about themselves.

The practical implication is not that HR should avoid uncomfortable initiatives around leadership teams with defensive tendencies; sometimes surfacing hard truths is exactly the job, and avoiding it entirely would be its own kind of failure. The implication is that the sequencing, framing, and support structure around an uncomfortable initiative needs to be calibrated to the actual psychology in the room, not the aspirational one. A leadership team with a genuinely defensive relationship to criticism may still need to hear hard feedback, but delivering it through a fully public, unfiltered channel, with no preparation and no private on ramp, is choosing the version of the tool least likely to produce anything other than a defensive crackdown. Reading Ego accurately isn't about avoiding hard truths. It's about choosing the delivery mechanism that gives a hard truth an actual chance of landing as information rather than as an attack.

CHAPTER 9: S FOR STAGE

An organization's developmental stage: startup, scaleup, mature, turnaround, decline shapes what almost every HR tool is actually for, which means the same tool can be doing entirely different jobs depending on when it's introduced, even when its mechanics never change.

Take structured interviewing, a genuinely well-evidenced hiring practice built around standardized questions, defined rubrics, and calibrated scoring meant to reduce bias and improve hiring accuracy. In an early-stage startup, hiring is usually happening fast, often through founder instinct and personal networks, introducing full structured interviewing too early can slow hiring to a crawl at exactly the moment speed matters most for survival; a startup that loses a critical hire to a slower, more rigorous process because a competitor moved faster has paid a real cost for a genuinely good practice introduced at the wrong stage. In a scaleup, hiring dozens of people a month with rapidly growing manager headcount; the same practice becomes close to essential. The informal, founder-instinct hiring that worked at fifteen people will produce wildly inconsistent quality and real bias exposure at two hundred; structured interviewing is precisely the tool needed in this stage. In a mature organization, the same practice risks calcifying into exactly the kind of box-checking bureaucracy it was originally designed to prevent, unless it's actively maintained and periodically revisited.

Turnaround and decline deserve separate mention, because they're the two stages most likely to be mistaken for each other despite requiring almost opposite approaches. A turnaround organization is under real pressure but has a credible, resourced path back to health, and initiatives there can reasonably ask for a degree of shared sacrifice and forward-looking optimism, because there's a believable story about what's on the other side of the discomfort. A declining organization, honestly assessed, doesn't have that credible path, and initiatives that borrow turnaround-style optimistic framing; "we're all in this together, better days ahead" - when the organization is actually in decline tend to land as either naive or dishonest to a workforce that can generally read the difference between real momentum and manufactured optimism, even when leadership itself hasn't fully admitted the difference internally.

The diagnostic question at the center of this dimension is simple to ask and often uncomfortable to answer honestly: what stage is this organization actually in, as distinct from what stage its leadership would like to believe it's in? Stage self-assessment is one of the more common blind spots at the leadership level; a mature company's leadership team, having spent years internalizing a scrappy, fast-moving self-image from its earlier years, will sometimes keep designing HR initiatives for a startup that no longer exists, because that's the version of the company they remember being. A declining company's leadership, for entirely human and understandable reasons, will sometimes describe their situation using turnaround language because turnaround is a more survivable story to tell investors, employees, and themselves than decline is.

Reading stage accurately, and separately from how leadership describes it, often means looking past the narrative and at more durable signals: headcount growth or contraction over the past several quarters, the actual tenure distribution of the workforce, how decisions get made now compared to two years ago, whether the organization is spending its energy expanding or defending what it already has. None of these require a leadership interview to assess, they're visible in the data any HR function already has access to, which makes Stage, in a useful way, one of the more empirically checkable dimensions in the whole framework, less dependent on reading a room and more dependent on simply being willing to look honestly at numbers a leadership narrative might prefer to obscure.

With all five lenses now on the table: Power, Ambient culture, Conditions, Ego, Stage; the framework is complete, but it isn't yet proven. A diagnostic tool only earns its place in a practitioner's actual practice once it's been tested against real decisions, in real organizations, where getting the read wrong has real costs. That's the work of Part III: the same HR tools, run through the same five lenses, landing in wildly different places depending on what those lenses reveal.

PART III - CONTEXT IN PRACTICE

Each chapter in this section follows the same structure on purpose. A single HR tool described plainly enough that any practitioner would recognize it. Two organizations, told as composites in the spirit of the note at the front of this book. The same tool, applied faithfully in both, producing close to opposite results. And at the end of each chapter, a short device worth explaining before it appears the first time: Two Lenses, one section written for the practitioner running the initiative day to day, one for the strategist deciding whether to greenlight it in the first place. The tool is identical in both contexts. The lens each reader needs to apply is not.

CHAPTER 10: HIRING & SELECTION

Structured interviewing is one of the most well-evidenced practices in the entire HR toolkit. Standardized questions asked of every candidate, a defined rubric scored consistently, interviewers calibrated against each other before they're let loose on real candidates; the research behind it, going back decades, is about as solid as anything in this field gets. It reduces bias. It improves prediction of on-the-job performance over unstructured, gut-feel interviewing. On paper, there is very little reason not to use it.

At a software company two years past its Series B in Yaba, growing from sixty to two hundred and forty people in eighteen months, structured interviewing arrived at almost exactly the right moment. Hiring had been running through a loose, founder-adjacent network: friends of early employees, warm intros, interviews that amounted to a good conversation and a gut check. That approach had worked when the company was small enough for the founder to personally vouch for nearly every hire. It stopped working somewhere around employee one hundred, when hiring managers who'd never met the founder's old college roommate started making wildly inconsistent calls, when two people hired for the same role six weeks apart turned out to have been evaluated on completely different criteria, when a pattern of who got hired and who didn't started raising real, warranted bias concerns. The head of talent introduced structured interviewing not as an abstract best practice but as a direct answer to a problem everyone in the leadership team could already see and was already anxious about. Hiring managers, several of whom had personally felt the sting of a bad hire slipping through gut instinct, welcomed a rubric that gave them cover and consistency. Within two quarters, interview-to-offer ratios improved, and more importantly for buy-in; a couple of visibly strong hires came through candidates who would have been screened out informally under the old system for not fitting the founder's mental picture of who belonged there.

At a boutique design agency off Awolowo Way, eleven years old, forty-two employees, run by a founder who had built the entire creative reputation of the firm on her own eye for talent, the identical structured interviewing process; same rubric philosophy, same calibration training, brought in by a newly hired Head of People eager to modernize the function landed very differently. The founder had built her hiring approach around a genuinely real skill: an instinct for creative talent that had, by any honest measure, produced an unusually strong roster of designers over a decade, largely through interviews that looked far more like portfolio conversations and gut chemistry checks than anything a rubric could fully capture. When the new structured process required her to score candidates against predetermined criteria and sit through calibration sessions justifying her scores against a rubric, she experienced it; accurately, from her own frame as an implicit accusation that her decade of proven judgment couldn't be trusted. She didn't openly resist the new process. She simply started making the real hiring decisions in informal conversations before the structured interviews even happened, reducing the official process to theater everyone could see through, which did more damage to trust in the HR function than the old inconsistent system ever had.

Run both organizations through PACES and the divergence stops being surprising. On Power, the software company was scaling past the point where founder-adjacent authority alone could scale with it, a broader coalition of hiring managers genuinely needed and wanted a shared standard. The design agency's power structure hadn't moved an inch; it was still, entirely legitimately, one founder's judgment call, and a process that asked her to subordinate that judgment to a rubric was asking her to give up power that had never actually been up for negotiation. On Ambient culture, the software company's workforce had recent, painful memories of inconsistent hiring; the new process arrived as a relief, not an imposition. The design agency's ambient culture ran the opposite direction: its entire identity, told and retold in how the company talked about itself, was built on the founder's uncanny eye, and a process that implicitly treated that eye as unreliable cut against the story the whole organization talked about why it was good at what it did. On Stage, the software company was exactly at the inflection point structured interviewing exists to solve. The design agency, small and stable, was nowhere near that inflection point at all.

The lesson isn't that structured interviewing is right for scaleups and wrong for founder-led creative firms as a rule. It's that the tool's value depends entirely on whether the organization has actually outgrown the judgment it's meant to replace - and that's a PACES question, not a question the tool itself can answer.

Two Lenses

Practitioner Lens: Before introducing structured interviewing anywhere, find out whose judgment the current hiring process actually runs on, and ask directly whether that person experiences hiring inconsistency as a problem they want solved, or as the exercise of a skill they take real pride in. The same rubric conversation lands as relief in the first case and as an accusation in the second, and the difference is almost entirely in how it's introduced.

Strategist Lens: Before approving a hiring process overhaul, check where the organization actually sits on Stage and Power; has hiring genuinely outgrown a single person's judgment, or is HR importing a scaleup solution into an organization that hasn't reached scaleup problems yet. Greenlighting the tool before the organization needs it spends trust for a fix nobody was asking for.

CHAPTER 11: PERFORMANCE MANAGEMENT

Stack ranking: forcing a distribution of performance ratings across a team, so that even a strong team has a meaningful bottom tier is typically the source of the most pointed disagreement in performance management. Its defenders argue it counters rating inflation and forces honest differentiation; its critics argue it manufactures internal competition and punishes strong teams for the crime of being strong. Both are right, depending on where it's used.

At a sales-driven scaleup selling enterprise software, stack ranking fit the organization's existing logic almost seamlessly. Sales itself already ran on a visible, numeric, comparative system; everyone could see everyone else's quota attainment on a leaderboard updated in real time, and the entire culture had self-selected for people who found that motivating rather than corrosive. Extending a forced distribution into performance ratings wasn't a foreign import; it was a natural extension of a comparative logic the whole revenue org already lived inside daily. Reps who landed in the bottom tier in a given cycle mostly understood why, because the same information had been visible to them all quarter on the leaderboard, and the rating rarely arrived as a surprise.

At a mission-driven nonprofit running youth mentorship programs, thirty-five staff, deeply collaborative by both necessity and identity, the same forced distribution introduced by a new operations director who had seen it work at a previous for-profit employer produced a very different chemistry. Program staff worked in tightly interdependent teams where a mentee's outcome depended on collaboration across three or four people's efforts, not on any one person's individually measurable output. Forcing a bottom tier onto that structure meant, almost by mathematical necessity, that someone on every high-functioning team had to be rated below their genuinely collaborative peers for reasons that had more to do with the forced curve than with any real gap in contribution. Staff who had built their professional identity around the organization's mission, not around individual competition, experienced the low ratings as a rebuke to values the organization itself had spent years telling them to hold. Two of the most tenured, most respected staff members left within the year, both citing the new ratings system by name in their exit conversations; not because their ratings were factually wrong on some technical level, but because being told to see a valued colleague as competition rather than a collaborator cut against everything the organization had recruited them to believe about the work.

The PACES read makes the divergence legible rather than mysterious. Ambient culture is doing most of the work here: the scaleup's ambient culture already normalized visible, comparative evaluation; the nonprofit's ambient culture was built, deliberately and for good reason, around collaborative identity, and stack ranking asked staff to betray the very value system the organization had recruited them into. Ego matters too, in a subtler way; the nonprofit's leadership genuinely believed in its collaborative, values-driven self-image, which meant a comparative rating system created a dissonance leadership itself hadn't fully registered until the exit interviews started arriving.

Two Lenses

Practitioner Lens: Before implementing any forced distribution, look for whether comparative evaluation already exists informally somewhere in the organization's daily life; a leaderboard, a visible metric, a known pecking order. If it doesn't exist anywhere yet, introducing it formally through ratings will land as a foreign imposition, not a formalization of something people already live with.

Strategist Lens: Before approving a performance system wholesale from a previous employer, ask what that system was actually built to solve there, and whether this organization's actual operating model, individually attributable output versus genuinely interdependent teamwork resembles the one it came from closely enough to bear the transplant.

CHAPTER 12: COMPENSATION & PAY TRANSPARENCY

Radical pay transparency; publishing full salary bands, sometimes individual salaries, openly inside an organization has become one of the most discussed compensation practices of the last decade, championed as both a fairness intervention and a trust-building exercise. Its outcomes, in practice, diverge more sharply based on context than almost any other tool in this book.

At a fast-growing climate tech startup with a genuinely young, flat, low-politics culture, publishing full salary bands by role and level did roughly what its advocates promise. Employees could see exactly where they sat and what growth looked like; managers had a shared, defensible reference point instead of ad hoc negotiation; a few real pay inequities, mostly the unintentional byproduct of inconsistent early-stage offers, surfaced and were corrected without much drama, because the organization's flat structure and relatively short history meant there wasn't much accumulated grievance for the newly visible numbers to attach themselves to.

At a decades-old regional healthcare system, recently merged from two formerly competing hospital networks with a long, visible history of internal factionalism between the two legacy staffs, the identical transparency initiative; full salary band publication, championed by a new CHRO who had seen it succeed at her previous tech-sector employer landed as an accelerant rather than a solvent. The organization already had a politically fractured culture, with each legacy network harbouring long-standing, partly justified suspicions that the other had been favoured in the merger's staffing and leadership decisions. Publishing salary bands didn't create new information so much as hand a factionalized workforce a precise, numeric weapon to confirm what each side already suspected about the other. Grievances that had previously been vague: "they always get the better roles" became specific and citable and pay conversations that might have taken years to surface organically all arrived at once, with no existing mechanism in place to process that volume of legitimate but suddenly urgent conflict. The CHRO's intent was trust-building. The actual effect, in that specific political context, was to arm an already-fractured workforce with ammunition faster than the organization's conflict-resolution capacity could absorb.

Power and Ambient culture are the dominant lenses here. On Power, the healthcare system's post-merger leadership structure was itself unsettled and contested, with neither legacy network fully trusting the other's representatives; introducing radical transparency into a genuinely unresolved power struggle gave every existing faction a new, harder-edged tool to fight with. On Ambient culture, the climate tech startup had comparatively little accumulated grievance for transparency to activate; the healthcare system had years of it, actively suppressed rather than resolved, and transparency didn't create the fracture, it simply stopped letting the fracture stay quiet.

Two Lenses

Practitioner Lens: Before rolling out pay transparency, map existing grievance, not just existing pay gaps. An organization with real but relatively fresh, low-politics inequities can absorb transparency quickly. An organization with old, unresolved factional history needs a conflict-resolution capacity built alongside the transparency, not after it.

Strategist Lens: Before approving radical transparency as a trust-building initiative, ask honestly whether the organization's current power structure is settled enough to survive the conflict transparency will surface, or whether it's importing a tool built for organizations with far less unresolved history than this one carries.

CHAPTER 13: CULTURE CHANGE & VALUES WORK

Values workshops; the familiar exercise of gathering leadership, sometimes the whole company, to define or redefine a shared set of stated values, usually producing a handful of words on a wall and a slide deck everyone half-remembers a year later are one of the most frequently run and most frequently mocked HR interventions in existence. The mockery is often earned. It's also, sometimes, unfair, because the same exercise can genuinely reset an organization's shared language, or can become an expensive, cynicism-generating exercise in saying things nobody actually intends to live by.

At a mid-sized logistics company emerging from a genuine turnaround; new CEO, restructured leadership team, a credible eighteen-month story of real operational improvement already visible in the numbers, a values workshop convened to articulate what the "new" company stood for landed as a moment of real closure. Staff had lived through the hard part already; the numbers backing up the turnaround were visible and undeniable before the workshop even happened, and the exercise functioned less as aspiration and more as ratification of a change people could already feel. The values that emerged stuck, not because the workshop itself was unusually well-facilitated, but because the organization had already earned the right to state them.

At a mid-sized retailer three years into a slow, unacknowledged decline; falling market share, two rounds of quiet layoffs described internally as "restructuring," a leadership team that had not publicly admitted the trajectory it was actually on, an identical values workshop, run by an outside consultant using nearly the same format, produced open cynicism before it even concluded. Employees who had lived through two rounds of layoffs that contradicted the company's stated value of "we take care of our people" sat through a session generating new language for values the organization had already visibly failed to honour, and several staff said so, directly, in the room. The workshop's content wasn't the problem. It was staged onto a Stage the organization hadn't honestly acknowledged, using Ego-driven language "we take care of our people" that the organization's recent actual behaviour had already disproven in the room's collective memory.

Two Lenses

Practitioner Lens: Before running a values workshop, ask what the organization's recent actions have already communicated about its real values, whether anyone intends to say so out loud or not. A workshop can ratify a genuine shift. It cannot talk an organization out of a pattern its own workforce has already observed.

Strategist Lens: Before commissioning culture work, be honest about Stage; a values reset works as closure after real change has already happened and works far worse as a substitute for change that hasn't happened yet.

CHAPTER 14: LAYOFFS & DOWNSIZING

Nowhere in HR practice does context matter more, and nowhere does an identical process produce a wider range of outcomes, than in layoffs. The mechanics: severance formula, communication timeline, manager talking points, outplacement support can be lifted almost word for word between two organizations and land as either a moment that preserves years of accumulated trust, or one that destroys it permanently, based almost entirely on precedent and timing rather than on the generosity of the package itself.

A mid-stage software company facing its first-ever layoff, a genuine reduction tied to a real strategic pivot rather than mismanagement, communicated it with unusual care: leadership explained the reasoning honestly rather than in euphemism, severance exceeded legal minimums by a meaningful margin, managers were briefed thoroughly enough to answer real questions rather than reading a script, and critically, this was the first time anything like this had happened at the company, so there was no accumulated cynicism from past broken promises to colour how the news landed. Remaining employees, watching how departing colleagues were treated, largely reported feeling that if it had to happen, it happened about as well as it could have. Retention and morale among the remaining team, while shaken, recovered within a couple of quarters.

A legacy manufacturing company running its fourth layoff in six years, each time described to the remaining workforce as "the last one" each time followed within eighteen months by another round, faced an almost mechanically identical situation; same severance formula, arguably more generous than the software company's, same careful manager briefings, same honest communication about the strategic reasoning. None of it mattered nearly as much as the fourth in the phrase "fourth layoff in six years". Remaining employees didn't evaluate this round on its own terms; they evaluated it against three prior rounds of the exact same reassurance that this would be the last one, and the reassurance itself had become a kind of dark joke on the floor, repeated with visible irony the moment leadership used the phrase again. Trust in leadership's word; not in the severance package, not in the process, but specifically in anything leadership said about the future had been spent down to almost nothing, and no amount of care in this particular round could rebuild what four rounds of broken promises had already cost.

This is Ambient culture and precedent operating at their most consequential, exactly as previewed in Chapter 6. The tool: severance formula, communication plan, manager training was not meaningfully different between the two companies. What differed entirely was what each workforce already believed about whether leadership's words could be trusted, built entirely from what had actually happened the previous times something like this was said. A first layoff, honestly handled, is remembered as a hard but fair moment. A fourth layoff, however honestly handled in isolation, is received against the memory of the first three and no communication plan, however well-crafted, can out-argue a workforce's own lived experience of what happened last time.

Two Lenses

Practitioner Lens: Before drafting layoff communications, find out explicitly how many times this has happened before, and what language leadership used each time; especially any promise that "this is the last one". Reusing that exact phrase, even sincerely meant this time, in an organization where it's already been broken once, will read as either naivety or dishonesty regardless of intent.

Strategist Lens: Before approving layoff communications built around reassurance about the future, be honest about whether this organization has the standing to make that reassurance credible. An organization on its first layoff can promise stability and be believed. An organization on its fourth cannot, and pretending otherwise costs more trust than admitting the uncertainty honestly would.

CHAPTER 15: DEI PROGRAMS

Few areas of HR practice have been more thoroughly commoditized into portable, off-the-shelf programming than DEI; unconscious bias training modules, employee resource group templates, representation dashboards, sold and delivered in nearly identical form across companies with wildly different politics, histories, and readiness to receive them. Few areas also demonstrate more clearly how identical programming, contextualized differently, can either build genuine goodwill or generate exactly the backlash its critics predict.

A professional services firm, already several years into deliberate, incremental work on inclusive hiring practices, introduced a formal DEI function with executive sponsorship, transparent representation goals tied to existing business objectives rather than treated as a separate initiative, and programming built collaboratively with input from employee resource groups that already existed informally before the formal program launched. The work landed as a natural extension of something the organization had already been visibly doing, not an abrupt import; because the firm had spent real time on Ambient culture and Power before formalizing anything, the formal program had constituencies already invested in its success before it launched, and leadership had already demonstrated enough genuine commitment, in Ego terms, that the program read as authentic follow-through rather than performative messaging.

A manufacturing company with a workforce split fairly evenly across a politically and culturally heterogeneous region, no prior history of any similar programming, and a leadership team that greenlit an off-the-shelf DEI initiative largely in reaction to industry pressure and a competitor's public announcement, rather than from any internally generated conviction, introduced a nearly identical program: same training modules, same representation dashboard, same resource group structure, purchased directly from a vendor and rolled out with minimal internal adaptation. Without any prior internal conversation preparing the ground, without leadership genuinely able to articulate why the program mattered to this organization specifically beyond "our competitors are doing it," and without any accounting for the workforce's actual range of views on the subject, the rollout was received by a significant portion of employees as an imposed, externally-driven initiative disconnected from anything they recognized about their own workplace. The backlash that followed wasn't a referendum on the underlying goals of the program; many employees who pushed back said, when asked directly, that they didn't necessarily disagree with the goals themselves. It was a reaction to how disconnected the program felt from any context the workforce actually recognized: no visible internal history building toward it, no leadership able to speak to it in their own words rather than the vendor's, and a launch that read as reactive box-checking rather than considered organizational commitment.

This is a chapter where every PACES dimension is doing real work simultaneously, which is part of why DEI programming is so often mishandled: it's rarely a single-dimension failure. Ambient culture determines whether the program builds on existing internal momentum or lands as a cold import. Ego determines whether leadership can speak to it with genuine conviction or only in borrowed vendor language, and employees are generally accurate readers of that difference. Conditions matter enormously too; a program launched reactively, in response to competitor or media pressure, carries a different, more fragile legitimacy than one launched from sustained internal conviction, and workforces tend to sense the difference even when leadership hopes they won't.

None of this is an argument against DEI programming, and it should not be read as one; it's an argument, consistent with the whole of this book, against the copy-pasted, vendor-templated version of it that skips the contextual groundwork entirely. The programs that succeed durably are, almost without exception, the ones built with real reference to the specific organization's history, politics, and leadership conviction; not the ones assembled fastest from the most polished template.

Two Lenses

Practitioner Lens: Before launching DEI programming, find out honestly whether there's already internal momentum to build on: existing informal resource groups, prior smaller initiatives, employees already doing some version of this work unofficially and build the formal program as an extension of that, rather than introducing it as something entirely new landing from outside.

Strategist Lens: Before greenlighting DEI programming, be honest about the actual motivation driving the decision, and whether leadership can articulate that motivation in their own words rather than a vendor's or a competitor's. A program launched from genuine internal conviction, and one launched from reactive pressure will use identical materials and produce very different receptions.

CHAPTER 16: M&A & INTEGRATION

Post-merger culture integration; the work of blending two previously separate organizations' processes, systems, and cultures into one is arguably the single hardest test of contextual competence in this entire book, because it requires reading not one organization's PACES profile but two simultaneously, and reconciling wherever they conflict.

A software company acquiring a smaller, earlier-stage competitor approached integration with a clear, publicly stated intention: the acquired company's product and technical talent were the entire rationale for the deal, and leadership deliberately preserved much of the acquired team's existing culture, tools, and even reporting structure for the first year rather than immediately absorbing them into the parent company's systems. This was, in PACES terms, an accurate read of both organizations' Power and Ambient culture; the acquired team's culture was precisely the asset being purchased and dismantling it immediately in the name of standardization would have destroyed the thing the deal was actually for. Integration proceeded gradually over eighteen months, with the acquired team's own leaders retained and given real authority, and the acquired company's staff largely stayed through the transition, because the acquisition, in practice, behaved consistently with what had actually been promised at the time of the deal.

A traditional industrial company acquiring a similarly sized regional competitor, motivated primarily by cost synergies and market consolidation rather than by any particular admiration for the acquired company's culture or people, ran an almost identical playbook of "preserve the acquired culture for a transition period", largely because that language had appeared successfully in a case study at an industry conference despite the underlying rationale for the deal being entirely different. The acquired company's staff, correctly reading between the lines of the deal's actual economic logic, understood fairly quickly that "preserve the culture for now" really meant "we haven't finished deciding what to cut yet", and the promised transition period became a slow-motion, low-trust waiting game rather than the genuine cultural respect it had been in the software company's case. When redundancies did eventually arrive, roughly on schedule with what most of the acquired staff had already assumed, the earlier language about culture preservation was remembered less as a genuine intention and more as a euphemism that had bought the acquiring company a year of quiet before the real plan revealed itself.

The difference wasn't the integration language, which was nearly identical in both cases. It was Conditions; the actual economic logic of each deal and whether leadership's stated Ego matched that logic honestly. The software company's stated respect for the acquired culture was consistent with the deal's actual rationale, and behaviour over eighteen months confirmed it. The industrial company's stated respect was inconsistent with a deal that was, from its inception, primarily about cost synergy, and the acquired staff's accurate read of that inconsistency shaped how every subsequent communication was received, regardless of how carefully worded it was.

Two Lenses

Practitioner Lens: Before communicating integration plans to an acquired workforce, understand the deal's actual economic rationale, not just the announced narrative, because the acquired staff will generally figure out the real rationale within the first few months regardless of what's said publicly, and every subsequent communication will be read against that discovery.

Strategist Lens: Before approving integration messaging, make sure the stated intention toward the acquired culture is actually consistent with the deal's underlying economic logic. A promise to preserve culture that contradicts the deal's real purpose will eventually be exposed by events, and the exposure costs more trust than an honest, less generous message would have cost upfront.

PART IV - BECOMING CONTEXT-LITERATE

CHAPTER 17: THE WEEKLY HABIT

A framework that only gets used once, during a formal planning phase for a major initiative, will save an organization from its biggest, most visible mistakes and do almost nothing about the dozens of smaller ones that accumulate quietly in between. Most of the damage documented in this book didn't come from a single catastrophic misread of a major rollout. It came from an accumulation of smaller decisions; a comment in a meeting, a policy tweak, an offhand promise from a leader, each one too minor to warrant a formal PACES review, each one still landing on a room that had its own power structure, ambient culture, conditions, ego, and stage, and each one still capable of quietly spending trust nobody was tracking.

The fix isn't a bigger process. It's a smaller one, run more often. Somewhere between the full diagnostic effort a major initiative deserves and the zero diagnostic effort most day-to-day decisions actually get, there's a lightweight version of PACES that takes minutes rather than weeks, and the practitioners who build it into an actual weekly habit tend to develop a kind of contextual reflex that eventually stops feeling like a checklist at all and starts feeling like the way they naturally think.

The habit itself is simple enough to describe in a single paragraph, which is the point. Once a week, before the busiest day on the calendar rather than after it, spend ten minutes reviewing whatever's coming up in the next seven days that touches people decisions in any way: a policy announcement, a difficult conversation, a new process rollout, even a routine meeting where something sensitive might come up and run each one through an abbreviated version of the five lenses. Not the full diagnostic interviews described in Part II. A faster gut-check: who actually has to be on board with this for it to stick, and have they actually said so. What's the history here, has anything like this been tried before, and how did it go. What's happening around this organization right now, externally, that might change how this lands. Is there a gap between how leadership describes itself and how leadership is likely to actually react if this goes sideways. And is this organization, or this specific team, in a different place developmentally than the plan currently assumes.

Ten minutes, five questions, run against everything on the calendar for the week; this is a small enough investment that it survives contact with a busy week in a way a more elaborate process wouldn't. The value isn't in any single answer generated during those ten minutes. It's in the repetition. A muscle exercised once a month stays a conscious, effortful exercise. A muscle exercised weekly, for long enough, becomes something closer to instinct and instinct, running in the background during the fast-moving meetings and hallway conversations where most real HR judgment actually happens, is worth more than any amount of diagnostic rigor applied only during formal planning phases.

There's a specific failure mode worth naming, because it's the one most likely to derail the habit before it takes root: turning the weekly review into a compliance exercise rather than a genuine diagnostic one. It's tempting, especially for practitioners who like structure, to build an elaborate template, fill in every field dutifully, and mistake the completion of the form for the completion of the thinking. The habit only works if the ten minutes are spent actually asking the questions honestly, including the uncomfortable answers - yes, this needs to be delayed; yes, that person hasn't really signed off; yes, we've done something like this before and it went badly and nobody's mentioned it. A weekly ritual that never once produces an uncomfortable answer isn't a diagnostic habit. It's theater, and it's worth being suspicious of your own practice if ten straight weeks of review never once flags anything worth changing.

Building this into an actual sustained habit, rather than a good intention abandoned by week three, benefits from a small amount of external structure. Some practitioners keep a simple running note; five questions, one line each, revisited every Monday morning that takes less time to maintain than it takes to describe. Others build it into an existing ritual they already have, attaching it to a weekly one-on-one with their own manager or a standing team meeting, so the habit inherits an existing structure rather than requiring a new one to be built from scratch. The specific mechanism matters far less than the consistency, and the consistency matters far less for what it produces in any given week than for what it trains, gradually, in the practitioner running it.

The deeper aim of this habit, beyond catching any individual mistake before it happens, is to close the gap this whole book has been describing between technical competence and contextual competence. Technical competence gets built through repetition too; nobody runs a single comp analysis and calls themselves fluent in compensation design; fluency comes from doing it enough times that the mechanics stop requiring conscious effort. Contextual competence works the same way, and yet almost nothing in standard HR training treats it as a skill that benefits from repetition, because almost nothing in standard HR training frames it as a skill at all rather than an innate sensibility some practitioners happen to have. The weekly habit is, in a real sense, contextual competence's equivalent of running the numbers on a comp study for the hundredth time: unglamorous, easy to skip, and the actual mechanism by which a conscious framework turns into an unconscious reflex.

CHAPTER 18: READING THE ROOM IN REAL TIME

The weekly habit described in the last chapter operates on a timescale of days. This chapter operates on a timescale of minutes, because a great deal of the most consequential contextual reading a practitioner ever does happens live, inside a single meeting, with no time to run a five-question checklist before responding.

Meetings are where an organization's actual PACES profile is most visible and most frequently ignored, because most people in most meetings are focused on content: what's being proposed, what's being decided; rather than on the layer underneath it: who's actually speaking for whom, what's being said versus what's being meant, and what the room's reaction is actually communicating as distinct from what it appears to be communicating on the surface.

Start with the gap between what's said and what's meant, because it's the single most useful thing to train an ear for. A leader who says "I want to make sure we've really thought this through" in response to a proposal is, depending entirely on tone, cadence, and everything the room already knows about that leader's history, either offering genuine, constructive caution or delivering a polite, socially acceptable version of no. The words are nearly identical in both cases. The meaning is not, and a practitioner who takes the literal content of that sentence at face value; treating it as an invitation to bring back more analysis while the room actually heard a soft rejection, will spend real effort addressing a question nobody was actually asking, while the real objection goes unaddressed and eventually resurfaces somewhere less convenient.

Silence deserves particular attention, because it is one of the most information-dense signals available in any meeting and one of the most commonly misread. A room that goes quiet after a proposal can mean several very different things: genuine agreement with nothing left to add, confusion nobody wants to admit to in front of peers, disagreement that feels too risky to voice given who's in the room, or simple disengagement from people who've already decided the outcome doesn't depend on their input. These four silences look identical from the front of the room and require completely different responses, and the only reliable way to tell them apart in the moment is to already have a read on the room's Power structure and Ambient culture walking in; who typically speaks first when there's real disagreement, who tends to go quiet specifically when they disagree rather than when they're confused, whose silence has historically preceded a problem surfacing later rather than genuine buy-in.

Power dynamics inside a single meeting are often more legible than an org chart, for a practitioner who knows what to watch. Notice who looks toward whom before answering a question; that micro-glance toward a particular person before responding is frequently a more accurate map of real authority than any reporting line, revealing whose approval someone is unconsciously checking for before committing to an answer out loud. Notice who gets interrupted and who doesn't, and whether that pattern tracks with title or runs against it. Notice whose questions get answered directly and whose get gently redirected; redirection is often a more polite, harder-to-call-out version of dismissal, and a practitioner attuned to it can hear a real objection being deflected in real time, long before it turns into the kind of quiet resistance that only surfaces months later as an initiative mysteriously stalls.

None of this argues for turning every meeting into an exercise in suspicious over-analysis, which would be exhausting and, frankly, would make a practitioner worse at actually connecting with people rather than better. The aim is closer to what a skilled clinician develops after years of practice: a background awareness running alongside normal, genuine engagement, occasionally surfacing something worth paying closer attention to, rather than a constant, effortful state of active surveillance. This is, again, a skill that improves specifically with deliberate repetition; the practitioner who leaves every meeting for a few months asking themselves what actually happened underneath the stated agenda will develop a faster, more automatic version of that same reading within a surprisingly short time, in much the same way the weekly habit from the last chapter eventually stops feeling like a checklist and starts feeling like instinct.

A final, practical note: reading a room accurately is only half of what this skill is for. The other half is knowing when and how to act on the read without embarrassing anyone or overplaying a hand built on an inference rather than a certainty. Naming a dynamic too bluntly and too publicly; "I notice everyone keeps looking at David before answering" can do more damage than the dynamic itself, turning an unspoken pattern into an uncomfortable, explicit one before anyone was ready to have that conversation openly. The more useful move is usually quieter: following up privately with the person who went silent, adjusting the framing of a proposal after noticing a pattern of subtle deflection, or simply factoring the real read into a decision without narrating the read itself out loud. Reading the room well is a private skill practiced in service of a public outcome, not a performance meant to demonstrate how perceptive the practitioner is.

CHAPTER 19: WHEN TO BREAK YOUR OWN RULES

Every framework, taught with enough conviction, eventually calcifies into the very thing it was built to critique. PACES is not exempt from this risk, and a book arguing against rigid, universally-applied best practices would be a strange place to end without confronting the possibility that its own five-lens diagnostic could become exactly that: another checklist, run mechanically, mistaken for wisdom rather than treated as a starting point for it.

There are real situations where the full PACES diagnostic is the wrong tool and naming them plainly is the only honest way to close this argument out. The clearest case is genuine urgency: a harassment complaint requiring immediate action, a safety issue, a legal exposure with a hard deadline. In situations like these, spending real time on a careful five-lens diagnostic before acting isn't rigor. It's a failure to recognize that some decisions have a different kind of context entirely: one where the cost of delay outweighs the cost of an imperfect read, and where the ethical and legal obligation to act decisively overrides the general principle that slower, more diagnostic work usually produces better outcomes. A practitioner who runs a careful PACES review before responding to an active safety concern has badly misapplied the very framework meant to prevent misapplied tools.

A second, subtler case is when a practitioner has genuinely internalized the reflex the weekly habit and real-time reading in the last two chapters are meant to build. The whole architecture of PACES, run formally, exists to compensate for the fact that contextual reading doesn't come naturally to most practitioners without deliberate practice. A practitioner who has done that practice for long enough sometimes develops a genuine, well-earned intuition that outpaces the formal exercise; and insisting on running the full explicit diagnostic anyway, every time, out of loyalty to the process rather than trust in the judgment the process was meant to build, is its own kind of rigidity. This is a real risk worth naming honestly rather than glossing over: intuition earned through years of deliberate practice is different from the untrained gut instinct this book spent its first three parts warning against and conflating the two; treating all intuition as equally suspect regardless of how it was built, would throw away the actual payoff of everything the weekly habit was meant to produce.

There is also the matter of over-applying PACES defensively; using the diagnostic not to genuinely assess whether an initiative fits its context, but as a sophisticated-sounding justification for inaction, risk-aversion, or simply not wanting to do something uncomfortable. "The context isn't right" can be an honest, well-reasoned diagnostic conclusion, or it can be a more articulate version of an excuse a less framework-literate practitioner would have made anyway. The only real defence against this misuse is intellectual honesty with oneself, which no framework can enforce from the outside; a practitioner using PACES as cover for avoidance will eventually notice, if they're paying attention, that their diagnostic conclusions suspiciously always favour the path requiring the least personal discomfort, and that pattern itself is worth treating as a signal.

The deepest version of this chapter's argument, though, isn't about edge cases and exceptions. It's about what kind of authority a framework like this is entitled to claim in the first place. PACES was built, across this whole book, in direct opposition to the idea that any tool, however well-evidenced, deserves to be applied the same way regardless of the room it's landing in. Holding PACES itself to a lower standard than every other tool this book has scrutinized; treating it as the one framework exempt from the very question of fit that it exists to ask about everything else would be a genuine failure of the book's own logic, not just a minor inconsistency. The correct relationship to PACES is the same relationship this book has argued for with every other HR tool across the last two Parts: useful, often essential, and still subordinate to the practitioner's own honest judgment about whether, here, now, in this particular room, it's actually the right instrument to reach for.

That closing loop is what keeps this book from becoming exactly the kind of universal playbook it opened by criticizing. The five lenses are a genuinely useful default. They are not a law, and the practitioner who forgets that distinction has learned everything in this book except the thing it was actually trying to teach.

CHAPTER 20: A CLOSING CASE

Joyce Marsh had been Head of People at a hundred-and-ten-person agricultural technology company for four months when the CEO asked her to fix retention among the field engineering team; the people who spent most of their working lives out at farms and test sites rather than in the office, installing and maintaining the company's soil-sensor hardware, and who were leaving at nearly triple the rate of every other department.

The instinctive, technically sound response was sitting right there in every retention playbook she'd ever used: run an engagement survey, benchmark comp against the market, and roll out a revised career ladder with clearer advancement criteria for field roles that had historically felt like a dead end compared to the office-based engineering track. All three of those interventions had worked for her before. All three were, in isolation, reasonable diagnoses of a retention problem. She had, by this point in her career, learned not to reach for any of them first.

She spent the first two weeks doing something that produced no visible output at all, which required a real conversation with the CEO about why the fix wasn't already underway. She ran the Power question first: who actually had to sign off for anything to stick, and quickly learned that the answer wasn't the CEO, despite his sponsorship of the project, but the VP of Engineering, a former field engineer himself who had built the entire field team's reporting structure around his own two decades of experience in the role and who had quietly resisted the last HR-led initiative aimed at his team eighteen months earlier; a standardized performance review process that had been imposed without his input and that he had never fully implemented, running a shadow version of his own instead. Any fix that arrived without his genuine buy-in, however well-designed, was likely to meet the same fate.

She ran Ambient culture next, and the eighteen-month-old shadow performance review turned out to be the tip of something larger. Field engineers, it emerged over several informal conversations with the longest-tenured member of the team, a nine-year employee named Dairo who had trained nearly a third of the current roster, had watched two previous "we're going to fix this" initiatives arrive with real fanfare and quietly evaporate within a couple of quarters, both times without any acknowledgment from leadership that they had been abandoned. The prevailing expectation on the team, largely unspoken but completely consistent across everyone she talked to, was that whatever came next would follow the same pattern: initial attention, gradual neglect, no accountability for the neglect.

Conditions added a layer the CEO had not mentioned in the original ask. The company was nine months from a Series C raise the board was counting on, and the metric investors would scrutinize most closely was customer churn - which, it turned out, correlated almost exactly with field engineer turnover, since long-tenured field engineers held relationships with farm operators that newer hires simply hadn't built yet. The retention problem wasn't a nice-to-have people initiative sitting alongside the company's real priorities. It was, whether leadership had fully connected the dots or not, sitting directly on the critical path to the fundraise.

Ego surfaced something more delicate. The CEO's public self-image, reinforced in nearly every all-hands, was built around being an unusually people-first leader in a notoriously unsentimental industry but his actual behaviour under pressure, observed twice during Joyce's first four months, leaned toward impatience with anything that couldn't show results within a single quarter. A fix that required real patience, and that might show soft signals of progress before hard retention numbers moved, risked losing his sponsorship exactly when it needed it most, unless it was framed from the outset with clear, visible early markers he could point to.

Stage rounded out the picture: this wasn't a startup's chaotic early growing pains, and it wasn't a mature organization's slow bureaucratic drift. It was a company at exactly the inflection point where informal, founder-adjacent management of field roles had stopped scaling, and where the VP of Engineering's decades of hands-on credibility, genuinely earned, hadn't yet translated into the kind of formal structure a hundred-and-ten-person company actually needed.

The plan Joyce eventually proposed looked, on paper, almost nothing like the survey-benchmark-ladder package she'd initially considered, and almost everything like something built for this specific room. It started not with a company-wide initiative but with a working group of four field engineers, including Dairo, co-designing the new career ladder alongside the VP of Engineering rather than around him;

giving him real ownership over something that had previously been imposed on his team without him. It deliberately avoided the word "initiative" in early communications, given what that word had come to mean to a team that had watched two of them quietly die, and instead described the work as "fixing what Dairo and the team have been asking about for a while", which was, plainly, true. It built in a visible, monthly checkpoint the CEO could point to in leadership meetings; not final retention numbers, which would take two quarters to move meaningfully, but concrete milestones the working group had actually completed, giving his impatience something real to hold onto without needing to see the lagging metric move before he could stay convinced it was working. And it explicitly, publicly acknowledged the two previous failed attempts in the first team-wide conversation about the new approach, rather than pretending they hadn't happened; a five-minute admission that cost the initiative nothing and bought it more credibility than any part of the actual plan.

Eleven months later, field engineering turnover had fallen by more than half, and as it happened closely enough to align with the completed Series C that the correlation made it into the board deck. None of the individual components of Joyce's plan were technically novel. A co-designed career ladder, a working group, a monthly checkpoint, an honest acknowledgment of past failure; none of it would look unusual in any HR toolkit. What made it work wasn't the tools. It was that every one of them had been chosen, sequenced, and framed specifically for this room, by someone who had spent two weeks finding out what this room actually was before deciding what to build inside it.

This is, in the end, the whole argument of this book, worked through one more time in full: the tools were never the hard part. Reading the room, they were about to land in always was. Context isn't a preliminary step to get through on the way to the real work of HR.

It is the real work of HR. Everything else is vocabulary.

PACES DIAGNOSTIC TOOL

Access the tool below by clicking on the pictorial. You need to be a registered member of the website to gain access to it. It compliments this book.