

Submerged

Hidden Innovation in the Age of AI

Hugo Lérias

2026

Contents

Foreword	5
PART I Beneath the Surface.....	10
CHAPTER ONE The Choice, and the Fleet	11
CHAPTER TWO Battle-Tested	25
CHAPTER THREE Why Now	35
PART II Why the System Buries the Best Ideas.....	49
CHAPTER FOUR The Gate	50
CHAPTER FIVE The Big Factory	61
CHAPTER SIX The Increment That Pays	73
PART III Why It's Never Been Fixed	89
CHAPTER SEVEN The Graveyard of Good Programs.....	90
CHAPTER EIGHT Promoted for the Big Bet.....	108
CHAPTER NINE The Silence That Follows	127
PART IV Surface or Submerge	143
CHAPTER TEN The Revelation Threshold	144
CHAPTER ELEVEN Stop Hunting the Fleet	164
CHAPTER TWELVE The Dual Operating System	186
CHAPTER THIRTEEN Does AI Break the Pattern?.....	213
CONCLUSION What To Do About It.....	232
AFTERWORD A Note on How This Was Made.....	238
APPENDIX The Promises Ledger	241
ABOUT The Author, and the Companion Volume.....	244

Copyright

Submerged: Hidden Innovation in the Age of AI

Copyright © 2026 Hugo Lérias.

This work is licensed under a Creative Commons Attribution–NonCommercial–NoDerivatives 4.0 International licence (CC BY-NC-ND 4.0). You are free to share it — copy and redistribute it in any medium or format — provided you give appropriate credit, do not use it commercially, and do not distribute modified versions. The full licence text is at creativecommons.org/licenses/by-nc-nd/4.0/.

This book is free. It has never been for sale and is not intended to be. If you paid for a copy, you were overcharged. The current edition always lives at submerged.lerias.org.

Its companion volume, *Leapfrog: A Corporate Engineer's Field Guide to Delivering AI on the Cloud*, is also free, at leapfrog.lerias.org. *Leapfrog* is the how. This one is the why and the when.

Written outside working hours, on personal equipment and a personal account, under written approval from the author's employer. Every claim in it is traceable to published research or to the author's own reasoning. No

internal or confidential material from any employer has been used. Maria is a composite. She is narrated, never interviewed, and is not a source.

The draft notice on the preceding page applies to this file in full.

Foreword

Leapfrog was written for one person: the engineer inside a large company who can see what AI makes possible and cannot get anywhere near it.

That engineer is told constantly what AI can do — by management, by vendors, by the innovation department. They are shown frontier models of extraordinary sophistication and reminded that the company must deliver on the AI promise. Integrate it. Find the savings. Show the results.

The effect is not motivation. It is paralysis. Underneath the announcements sits a genuine fear of failing, and that fear has a history.

Anyone who worked in IT through the SaaS years watched the knowledge leave the building. It went to contractors, to consultancies, to vendors. What remained was a team held responsible for outcomes it no longer had the skills to produce. Now that same team is asked to deliver AI. No wonder nobody volunteers.

So the question worth asking is not whether AI is powerful. It is: what does an internal team actually need to know?

For a long time the answer has been the cloud. The move to cloud is one of the key enablers for approaching AI today. Not the only one — and not an invitation to buy everything the cloud

provider tells you that you need. The real gift of cloud was self-service: the ability to solve your own problem without asking permission first. AI makes that faster still. Put the two together and you have, in my opinion, a perfect storm.

Meanwhile the frontier model companies are busy defending their rankings on a benchmark. Latest, greatest, biggest. But we passed the point of sufficiency a while ago. The capability needed to do the simple things — the things that would actually change your company — has been sitting there for some time.

Take the invoice exercise in *Leapfrog*. Reviewing travel expenses is not a glamorous problem. Changing it so that nobody has to review them at all is a real one, and it is solvable today. How you implement it is your business. Should you buy a SaaS product for it? I don't think so. Many of these workflow tools already have AI built in, and some are open source. Why pay a monthly bill for something you can have out of the box?

That is what the first book is for. *Leapfrog* is the practical one — the book for the person with some cloud knowledge who wants to get on the AI train.

Submerged asks a different question. Why do we see so little innovation coming out of employees?

I believe innovation has to come from the people doing the work.

This is not a complaint about innovation departments. Their job is real, but it is a different job: scouting, watching what is happening outside, running the occasional proof of concept, cataloguing it, keeping other teams roughly aligned on what is on the radar. That has value. It is simply decoupled from the delivery of value. Nobody in that chain is the person with the problem.

Innovation needs to happen where it is used. The job of everyone else is to equip the people who use it — to put

genuinely powerful technology in their hands, safely, without requiring a budget cycle to test an idea.

Which brings me to the problem I have watched inside companies I have worked for, seen in companies I haven't, and then found waiting for me in the research.

Submerged innovation. Hidden innovation. Submarine innovation.

There are more submarine projects running in your company than you think. The question is why. There are several answers, but the first one is simple: people are afraid of failing.

Here is an unscientific way to put it. Some of us, walking into a shop, don't want to talk to anybody. We want to look around first, at our own speed, and work out what is there. Only when we have a question do we go and find someone. Innovation works the same way. A team has an idea. They don't yet know whether it is any good. They need room to find out before anyone asks them to justify it.

But fear is only the beginning. When somebody surfaces a new project, the politics of departments arrives with it. *Why did they do that? I had the budget for that. I didn't deliver it – and now they have.* Then comes the cycle of friction and destruction.

Destruction is the right word. These projects are not politely ignored. They are stopped. And the person who surfaced one and watched it die has learned something they will not unlearn: don't do that again.

So the loss is never one project. The loss is every idea that person will never bring you. It dies quietly, and you will never know it existed. What follows is worse than fear – it is futility. Fear is loud and gets the attention. Futility is quiet, and it is the one that does the damage.

I should include myself here. I have been building small tools with frontier models – some for work, some just for home. Nothing dramatic. And at some point the obvious thought

arrives: if I am doing this, and the colleagues sitting behind me are doing this, then there must already be an enormous amount of innovation built and hidden inside this company. A whole submarine fleet, waiting to be found.

But what if it never surfaces?

There is also the money. Companies think in cycles — budget planning, approval, allocation. Innovation does not run on that clock, and it never has.

And here is the part for the executives, which is the heart of this book. If you force the fleet up, if you go hunting for it, you will not find it. You will kill it. People will dive deeper. They will stop using company machines. They will build on their own phones, in their own accounts, at home, where your search and your controls cannot reach. You will not have sunk the submarine. You will have pushed it into international waters — and now you have lost the visibility, lost the control, and gained a data problem you cannot see.

Not everything needs to surface, either. Some of these tools are one person's helper, and that is fine. If it helps nobody else, let it be. The point was never that everything must come up. The point is that surfacing should be a door people are offered, not a net thrown over them.

So the answer is not to hunt. Give people a safe place to build underwater, at their own speed. Then build the thing that lets a good idea come up when it is ready.

That is what this book is about.

The research led me first to a concept called bootlegging, and from there to a question: has this always been happening? What does the evidence actually say — and is there anything in it for large manufacturing companies in particular, where the capital is heaviest and the doors only swing one way? How do we help them?

That was the work. I did it with AI, of course. I collected a great deal of research, argued with it, and built it through the chapters of this book, trying to get to a better position — and to get the fleet to surface.

Enjoy it, as always. There will be a companion site with the research behind the book, and beyond it, where I will keep offering my opinion and the opinions of others as we go.

That's it. Have fun.

PART I

Beneath the Surface

A book about hidden work has to begin by establishing that there is any, which is harder than it sounds, because the subject is defined by not being visible from wherever the reader happens to be sitting.

So Part I does one job, in three chapters. The first puts a person in front of you and then shows you how many of her there are. The second answers the objection that this is a story about artificial intelligence, by walking the same behaviour back through six decades and four technologies and showing that nothing about its shape has moved. The third asks what did change, because something plainly has.

There are no remedies in this part, and that is deliberate. Most of the fixes people reach for at this stage are aimed at a problem they have not yet finished looking at.

CHAPTER ONE

The Choice, and the Fleet

The meeting invitation arrived on a Tuesday, and Maria read it twice.

Quarterly review – AI use cases. Please come prepared to share what your team is working on.

She had something to share, and that was exactly the problem.

Eight months earlier she had been drowning in supplier deviation reports. Every non-conforming part that arrived at the plant generated one, with a description, a photograph, a disposition and a root-cause note, and her job was to read them, categorise them, spot the patterns and escalate anything systemic. On a good week there were forty of them. On a bad week there were two hundred, and by Friday evening the patterns she was paid to notice had simply disappeared into the volume.

So one evening, at home, she pasted an anonymised report into a chatbot and asked it to summarise the failure mode in one line. It did. She tried ten more, and it got them right – not perfectly, but better than she did at six on a Friday. Over the following weekends she built a rough workflow: strip the identifying data, batch the reports, extract the failure mode and

the supplier, cluster them, flag anything appearing three times in a fortnight.

It took her about six hours to build and it saved her roughly a day a week. After two months she caught a recurring seal defect five weeks earlier than the old process would have. She didn't tell anyone how she caught it. She just caught it.

Nobody asked her to build this and nobody approved it. It runs on her own account and she pays for it herself, eighteen euros a month, which she has never expensed, because expensing it would mean explaining it.

Now her manager wants to hear about AI use cases.

Maria has a good one, and she is sitting at her desk with the invitation open, doing the arithmetic that this whole book is about.

. . .

Here is what that arithmetic looks like from where she is sitting.

If she raises her hand, the tool becomes visible, and visibility means assessment. IT will want to know which model it uses, where the data goes, and whether an anonymisation script written by a quality engineer on a Sunday is sufficient control for supplier information. The honest answer is *probably, but I'm not certain*, and in a governance review "probably" is a no.

Legal will want to know whether an unapproved consumer account processing supplier data constitutes a policy breach. It does. She knows it does. She has known it did for eight months, and it is a thing she has tried not to think about too precisely.

And Quality Systems will hear about it. Quality Systems has a programme for exactly this, running for eighteen months now, with a budget she has heard described in terms she found hard to believe, involving two consultancies, and it has not yet shipped anything to her plant. She has sat through three of its update presentations. The roadmap slide is very good.

Maria's tool works today. Their programme will work eventually. Put those two facts side by side in a meeting and they do not read as *how useful*. They read as *how embarrassing*.

If she says nothing, none of that happens. She keeps her day a week, the seal defects keep getting caught early, and everyone carries on assuming Maria is simply good at her job, which she is, though not by the margin they think. The tool stays alive exactly as long as she does. When she moves to another plant next year, and she will, it goes with her, undocumented, and the reports start piling up again, and nobody will ever know that for two years this plant had a capability it has now lost, because the capability was never written down anywhere except in Maria's head and a folder on Maria's laptop.

That is the choice. Surface it and probably lose it, or hide it and definitely lose it, later, quietly.

I want to leave her there for now. We will come back to her, and by the end of the book we will know what she should do. At this point neither of us does.

. . .

The fleet

The temptation is to read Maria as an anecdote, but she is closer to a census.

In 2025, MIT's Project NANDA published a study of how generative AI was actually being used inside companies, as opposed to how it was being reported upward. The number that mattered was the gap between two others. Around 40 per cent of the companies studied had bought official large-language-model subscriptions. Around 90 per cent of employees were using personal AI tools for work.

Nine in ten people, using tools their employer has not sanctioned, to do the work their employer is paying for. That is

not a rebellious fringe in the engineering department. That is most of the building.

The same study found something more uncomfortable. Roughly 95 per cent of formal corporate GenAI pilots — the funded ones, the ones with a steering committee and a name — produced no measurable impact on profit and loss, while the unsanctioned personal use happening on private accounts at kitchen tables was delivering faster adoption and better returns.

The official fleet is in dry dock and the real one is already at sea, and nobody in the building has a map of it.

That 95 per cent figure gets over-read, including by people who should know better. The comparison is not clean. People using personal AI tools have self-selected for enthusiasm and are mostly doing easy, well-bounded tasks, while the formal pilots are attempting hard integration work under real compliance constraints. It is not a fair fight, and “shadow beats official” is not a conclusion the data can carry.

But a weaker conclusion survives, and it is more than enough for what follows: an enormous amount of real value is being created below the waterline, and the organisation cannot see any of it.

That is the submarine fleet. Maria is one boat.

. . . .

This is not new. It is only newly cheap.

The instinct, when you first see the shape of this, is to treat it as an AI story. AI is what made it visible to me, and it is certainly what makes it urgent now, but the pattern is older than the technology by about sixty years.

In 1963, Donald Schön published an article in the *Harvard Business Review* about how radical inventions actually got developed inside corporations. His answer annoyed people.

They got developed by champions who worked around the official process, often in secret, frequently in defiance of explicit instruction. He gave the behaviour a name that stuck: bootlegging.

The name has been reinvented roughly once a decade since. Skunkworks. Pet projects. Friday afternoon work. Creative deviance. Shadow IT. Stealth innovation. Now shadow AI. Each generation of managers discovers it, names it, and treats it as a novel symptom of whatever technology has just arrived, which is a bit like discovering rain.

The academic definition is worth having exactly, because every clause of it does work later on. Bootlegging, in the research literature, is a non-formalised, non-declared, bottom-up innovation process undertaken for the benefit of the bootlegger's own firm. Non-formalised: it isn't in the plan. Non-declared: nobody upstairs knows. Bottom-up: it starts with the person holding the problem. And for the benefit of the firm, which is the part that surprises people.

Because the natural assumption about hidden work is that hiding implies wrongdoing. Someone is slacking, or empire-building, or running a side business on company time. The research says otherwise, consistently and across decades. When Masoudnia and Szwejcowski interviewed 55 bootleggers across 34 high-technology firms, the dominant motivation they found was not personal gain at all. It was to benefit the organisation. The secrecy was never the point. The secrecy was the workaround.

Maria is not stealing from her employer. She is subsidising her employer at eighteen euros a month, and hiding the receipt.

There is a distinction in this literature that will run through the whole book. Permitted bootlegging is the sanctioned kind: 3M's fifteen per cent, Google's twenty per cent, the officially blessed sandbox. Hardcore bootlegging is the other kind, work

that continues after somebody with authority has said stop, and researchers have a name for that one too, creative deviance. We will spend real time on it, because the distance between “we allow this” and “we allow this until it becomes inconvenient” turns out to be where most of the damage happens.

. . .

The part that should worry you

If the story were simply that employees do unauthorised work and it is mostly fine, this would be a short book with a sensible conclusion about governance.

The complication is that the hidden work is not merely tolerable. It appears to be better.

The most rigorous study we have is by Paola Criscuolo, Ammon Salter and Anne ter Wal, published in *Organization Science* in 2014. They studied bootlegging inside a large R&D organisation quantitatively, at the individual level, looking at who did it, how much, and what came of it. Their finding was that bootleg effort was associated with achieving high levels of innovative performance. Roughly nine per cent more time spent underground raised the probability of a top innovation rating by about seventeen per cent.

Their explanation for *why* is the hinge the entire book turns on, so I want to state it plainly. Working underground gave people two things they could not get above the surface. It let them explore uncharted territory, directions nobody would have approved in advance because nobody could have justified them in advance. And it let them delay the assessment of embryonic ideas.

Delay the assessment. Not avoid it — delay it. An idea in its first weeks is fragile, half-wrong and impossible to defend, and it cannot survive a gate review, not because it is bad but because it

is young. So the submarine is less a hiding place from judgement than a nursery for things not yet ready to face it.

Which sets up the cruelty at the centre of the whole thing. The process that would kill the idea early is the same process that would fund it later. You cannot get to the version worth approving without a period in which it is not approvable.

Maria's tool was garbage for the first three weekends. If she had presented it in month one it would have been rejected, and rejected correctly. By month eight it works. There was no path from the first version to the second that ran through a meeting.

There is also a sting in the Criscuolo findings that we will come back to in Part II. The bootlegging advantage did not hold universally. When the organisation went through periods of formalising its R&D process, with more structure, more gates and more discipline, the people who increased their underground work became *less* likely to innovate rather than more. The submarine only works in certain waters. Tightening the process did not stop the hiding; it stopped the hiding from working.

. . . .

Why she can't just tell someone

Come back to Maria for a moment, because there is a piece of her arithmetic I skipped, and in my experience it is the piece most executives miss entirely.

The obvious question is why she doesn't simply raise it. Her manager is reasonable. The company says it wants exactly this. There is a slide about it.

The answer is that Maria is not weighing whether her tool is good. She is weighing whether it will be seen, and those are different problems with different mechanics.

Large organisations run on a scarce resource that nobody puts in the budget, which is attention. William Ocasio made this the centre of one of the most cited arguments in strategic management: firms behave the way they do because of how they channel and distribute the attention of their decision-makers. What gets acted on is what gets noticed. Merit does the filtering later, if at all. Notice comes first.

So how does attention get allocated? Cyril Bouquet and Julian Birkinshaw studied 283 subsidiaries of multinational companies to find out how units capture the attention of headquarters, and their answer had two parts. The first they called weight, meaning your structural position: how big you are, how significant your market, how much of the company's mass sits behind you. Weight pulls attention automatically and never has to ask for it. The second they called voice, meaning initiative-taking, profile-building, the active work of being noticed. Units without weight have to use voice, and it does work, but it costs something and it can be spent.

The property that makes this bite is that attention is zero-sum. Headquarters cannot attend to everything, so every increase in attention paid to one unit is a decrease in the attention available to everything else.

Now look at the Quality Systems programme through that lens. Eighteen months, two consultancies, a budget large enough to be quoted in disbelief. It is not competing with Maria for attention, it is the attention. It has mass. It does not need to be good, or fast, or shipped, to pull the room toward it. It only needs to be big, and it is.

Maria's tool has no weight at all. One person, six hours, eighteen euros a month.

So it will not be rejected, because rejection would require somebody to look at it first. It will simply never be seen, which is a quieter fate and one that leaves no record at all. Nobody files a

decision. There is nothing to appeal. The idea doesn't die, it just never registers.

This is the first of three scarce resources we will find bent in the same direction. Attention is one. Capital is another, and in a manufacturing company the capital story is brutal, because the decisions are one-way doors and the machinery guarding them was built for cathedral-sized bets rather than for six hours on a Sunday. Careers are the third, because the people promoted to allocate the first two are, systematically and measurably, the people most drawn to the big visible bet.

Three currencies, three separate mechanisms, three bodies of research that have never really been introduced to one another, and all of them pushing the same way.

That convergence is the argument of this book. Each of the three currencies is well documented on its own and I will cite the research as we go. What I think nobody has done is set them side by side and notice they all push the same way, and that connection is mine to make and mine to defend. You will not have to take it on faith. You will watch all three act on Maria long before the book asks you to agree with me about anything.

. . .

A curtain in a factory

I want to plant one thing before we go further, because you will need it much later and it is better if it is already sitting in your mind.

In 2012 a researcher named Ethan Bernstein published a study conducted inside the world's second-largest mobile phone factory. He was not looking for hidden innovation. He was looking at transparency, the idea, popular then and popular now, that the more managers can see of the work, the better the work gets.

What he found was that the workers on those lines had developed better methods than the official ones — faster and more reliable — and that they hid them. When managers walked the floor the line reverted to the procedure in the manual, and the improvements went back underwater until the coast was clear.

The suggestion boxes on every line were empty.

Then Bernstein did something almost comic in its simplicity. He hung a curtain.

I will leave it there. We will come back to what happened, and it will matter more once you have seen what the rest of the machinery does.

. . . .

What this book is going to do

Three questions, in order.

Is this permanent? If hidden innovation is a quirk of a particular era or a particular management fashion, then it is a problem you can wait out. If it recurs identically from Schön's 1963 laboratories to Maria's kitchen table, then it is wiring, and you have to design for it. That is the rest of Part I, and I think the answer will surprise you, because the age of this research usually gets treated as its weakness. In my opinion it is the opposite. A finding that survives sixty years, four technology revolutions and every management fad in between is not dated. It is load-bearing.

Why does the system do this? Not out of malice and not out of stupidity, which is the disappointing part. Parts II and III show a machine that suppresses its own best small ideas while behaving, at every individual step, entirely reasonably. The capital process is right to fear one-way doors. The gate is right to demand evidence. The executive is right to prioritise. Each

mechanism is defensible on its own and together they are catastrophic. And then the harder question, the one I think should have been asked decades ago: if all of this is so well documented, why has nobody fixed it? The answer to that is the darkest thing in this book.

What do you do? That is Part IV, from both seats — Maria’s, and the one across the table from her. I will tell you now that the answer is not “surface everything” and it is not “let a thousand submarines sail.” It is stranger than that, and it runs against the instinct of every executive I have ever watched encounter this problem for the first time. The instinct is to go and find the fleet. That instinct is the single most expensive mistake available here, and I will spend a chapter showing you exactly what it costs.

. . . .

Where the evidence stops

This chapter rests on four studies and they are not equally strong, so it is worth saying which is which before the rest of the book starts leaning on them.

Schön’s 1963 article is an argument rather than a dataset. Its authority comes from how often the finding has been rediscovered since, by people who were not reading Schön and mostly did not know he existed, which is a different kind of evidence and in some ways a better one, but it is not the kind you can put a confidence interval on. Masoudnia and Szejczewski interviewed fifty-five bootleggers in thirty-four firms and asked them why they did it. All thirty-four firms were high-technology, and a self-reported motive is the most flattering number a person ever produces about themselves. Criscuolo, Salter and ter Wal is the rigorous one, quantitative and at the individual level, and it is one organisation. The NANDA study is a year old and describes a technology that

changes faster than anyone can finish studying it, and I have already said what its headline figure cannot carry.

None of them is the finding. The finding is that four different methods, in four different decades, in industries with nothing much in common, all point the same way, and that after sixty years of people looking, nothing of comparable weight points the other way. That is what the chapter is claiming, and it is a weaker claim than any single one of those studies would let me make on its own.

The place it could break is the fleet itself. I have asserted that the submerged work is large — most of the building, not a fringe — and the only direct measurement of that in this chapter is a survey of self-reported tool use, which counts people who use a chatbot for work and not people who have built something.

So here is the study I would like somebody to run. Go into one organisation of any size, with an amnesty in writing, and count what people have actually built on their own accounts: how many things, how old, how many other people use them, and how many of the builders' own managers already knew. If the answer is that there are a handful, that they are trivial, and that the managers mostly knew, then there is no fleet and this book is about a rounding error. My expectation is that the last number is the one that will surprise people.

Maria illustrates the mechanism. The research carries it. Those are different jobs and I have not confused them.

. . .

Surface Check

So — *should Maria raise her hand on Tuesday?*

Not yet. Right now we know only that her tool is good and that she is not alone. We don't know what would actually happen to it in that room, or why a programme with eighteen months and two consultancies still hasn't shipped, or whether the room is even capable of seeing something that small. Those are the next chapters. Ask again then.

INSTRUMENT

Find Your Own Fleet

Before the next chapter, an exercise. It takes twenty minutes and it is uncomfortable, which is rather the point.

Do not send a survey. A survey is a request for a declaration, and declarations are exactly what this population has learned not to make. You will get zeros, you will believe them, and you will be wrong.

Instead, ask yourself these five questions and answer them honestly, on paper.

1. **What do I do in a week that I have quietly made easier, without telling anyone how?** Start with yourself. If you cannot find one, you are either unusually compliant or not looking hard enough.
2. **Who on my team is faster than their process should allow?** Not the person who works late. The person whose output doesn't reconcile with their hours. That gap is a boat.
3. **What was the last idea somebody brought me, and what happened to it?** Then the one before that. Go back five. This is the most important question in the exercise and Chapter 9 will explain why.

4. **If someone on my team were paying eighteen euros a month out of their own pocket to do their job better, how would I ever find out?** If the honest answer is that you wouldn't, you have just measured your visibility.
5. **What is the biggest programme in my area, and what has it shipped this year?** And then: what would it take for a six-hour tool to get five minutes in the same room?

You will notice that none of these questions asks anyone to confess anything, and that is deliberate. By the end of Chapter 11 you will understand exactly why asking directly is the one method guaranteed not to work.

Next: Chapter 2 — why almost all the research behind this chapter is old, and why that turns out to be the strongest thing about it rather than the weakest.

Maria is still sitting at her desk. The meeting is Tuesday.

CHAPTER TWO

Battle-Tested

You may have noticed, if you glanced at the sources behind the last chapter, that a lot of them are old.

The study that named the behaviour is from 1963. The factory curtain is from 2012. The work on how attention gets allocated is from the late nineties. For a book whose central character is running a large-language model on her laptop in a year that starts with a two, that is a strange bibliography.

Hold that thought, because by the end of this chapter I want it to have flipped completely, from the thing you would use to attack the argument into the thing that makes the argument hard to attack. The age of this evidence is not a weakness I am hoping you won't notice. In my opinion it is the whole point.

So let me show you the same behaviour surfacing, under a different name, in every decade since the Beatles released their first album, and you can watch what changes and what doesn't.

. . . .

1963: it gets a name

Donald Schön was studying how genuinely new inventions made it out of corporate research labs. Not incremental

improvements, but the real departures, the products that didn't fit the existing roadmap. He expected to find a process. What he found was people going around the process.

The interesting inventions, he reported, tended to have a *champion*: someone who believed in the idea and was willing to protect it, quietly, from the organisation that would otherwise have killed it in review. The champion often worked outside the official plan and sometimes in direct defiance of it. Schön gave the behaviour a name that survived, bootlegging, and made a claim that I think should have been far more disturbing to his readers than it apparently was — that for certain kinds of innovation the informal, hidden route was not a failure of the system but the only route that worked at all.

Note what is already present in 1963, before the personal computer, before the internet, before anything you would call technology in the modern sense. The hiding is there. So is the pro-organisational motive, because the champion is not stealing anything, he is protecting an idea from his own employer. And so is the structural cause, a formal process better at killing young ideas than at growing them.

Every element of Maria's Tuesday is sitting in Schön's 1963 paper. Only the tools have changed.

1970s–80s: it builds the airplane, and companies try to tame it

The next thing the pattern does is get famous.

Lockheed's Advanced Development Programs, the Skunk Works, is where most people first met the idea. Kelly Johnson's team produced the U-2 and the SR-71 Blackbird by operating deliberately outside the parent company's normal process: small, autonomous, secretive, insulated from the reviews and sign-offs that governed everything else. The lesson corporations took

from this was seductive and, I think, slightly wrong — that if you want radical output you build a walled garden and let a special team play in it.

I say slightly wrong because of what happened next, which is the more revealing move. Companies noticed that unofficial work produced value and they tried to domesticate it. If bootlegging works, the reasoning went, let us authorise it. 3M formalised its fifteen per cent time. Later, famously, Google offered twenty per cent time. The underground would be given a sanctioned room to happen in.

Hold on to that instinct, because we are going to watch it fail, repeatedly, for the rest of the book. For now just register that by the 1980s the pattern was strong enough that large, sophisticated companies were building official programmes to try to capture it. You do not build a programme to capture a fluke. You build one to capture something you have noticed keeps happening whether you like it or not.

1990s: it moves into software, and it gets a second name

When work moved onto computers the pattern moved with it, and picked up a label from a completely different discipline.

The information-systems world started writing about shadow IT, meaning systems, tools and workarounds that employees built or bought without the sanction of the IT department. The spreadsheet that quietly runs a business unit. The little database nobody approved. The tool the team depends on that appears on no official inventory. Silic and Back, studying it, described people bypassing sanctioned systems not out of malice but to close the gap between rigid corporate tooling and what the actual work demanded.

Same behaviour, same motive, new name, because a new profession had discovered it and assumed it was theirs.

This is the moment the pattern's disguise turns into a real problem. Each field renamed it, each field studied it in isolation, and nobody stood on a chair and said that this is the same thing Schön described in 1963. The R&D scholars had bootlegging. The IS scholars had shadow IT. The strategy scholars, as we will see, had their own vocabulary again. One phenomenon wearing a different coat in every room, so that nobody quite noticed how old and how general it was.

2000s–2010s: it gets caught, and it gets punished

By now the pattern has a rap sheet, and this decade adds the fact that matters most for Maria, which is consequences.

Researchers began looking not just at whether shadow work happened but at what happened to the people who did it. One experimental study, by Allen and colleagues, is worth stating precisely, because it is the empirical spine under Maria's fear. They looked at how employees who used shadow IT were evaluated, and found that those employees were less likely to receive a bonus in *both* the good-outcome and the bad-outcome conditions.

Read that twice. It did not matter whether the shadow tool worked. The penalty attached to the circumvention itself. You could break the rule and succeed, or break the rule and fail, and either way the evaluator marked you down for having broken the rule.

That is the exact calculation Maria ran in Chapter 1, and now you know it is not paranoia but a documented pattern. The organisation punishes the going-around rather than the outcome of the going-around, which means a rational employee learns to hide the going-around, which produces precisely the invisible fleet we started with.

2025: it goes mass

And then the tools got cheap, and the pattern stopped being a minority sport.

We saw the number in Chapter 1: roughly ninety per cent of employees using personal AI tools for work, against forty per cent of companies with official ones. What took a champion's courage in 1963, a skunkworks' budget in 1975 and a clever engineer's evenings in 1995 now takes eighteen euros and a Sunday afternoon. The barrier to running your own submarine collapsed, so the fleet grew from a few daring captains to very nearly everyone.

That is the whole change. Not the behaviour, but the cost of the behaviour. We will spend Chapter 3 on what that collapse actually means, because it is the reason this book is being written now rather than in any of the six previous decades when it could have been.

. . . .

The thing that didn't change

Line those decades up and look at what stayed constant while the technology underneath was replaced four times over.

The hiding stayed. The pro-organisational motive stayed, and in every era, across every rename, the dominant finding is that these people are trying to help their employer rather than rob it. The structural cause stayed, which is a formal process systematically better at stopping young ideas than at growing them. And the disguise stayed too, with each generation greeting the pattern as a novel symptom of its own new technology, naming it afresh, and missing that it was looking at something with a sixty-year pedigree.

Four technology revolutions and one unchanged pattern. A behaviour that survives that does not depend on the tools. It depends on something a great deal more stable than the tools, which is the shape of the large organisation itself, and it will therefore not be fixed by a better tool, a new policy, or the current transformation programme. It is wiring.

This is what I mean by battle-tested, and the phrase needs discipline, because it is very easy to abuse.

What “battle-tested” is allowed to mean

There is a lazy version of this argument that says the sources are old, therefore they are proven. That is not true, and I want no part of it. Old and wrong is a common combination, and age on its own certifies nothing at all.

The claim is narrower than that, and stronger. A finding is battle-tested when it has been re-tested — examined again, by different people, in different places, in different decades, under different conditions — and keeps surviving. It is the survival across re-tests that carries the weight, not the birth date.

Take the sharpest example, one we will lean on hard in Part II. In 1987, research found that companies set their investment hurdle rates, meaning the return a project must clear to get funded, far above their actual cost of capital, and then wait until an opportunity clears that inflated bar before committing. An old finding, and easy to wave away as a relic of eighties finance.

Except that it kept getting re-tested. Through the two-thousands. Through the twenty-tens. And into the twenty-twenties, with interest rates having collapsed and then whipsawed in between, the finding held: hurdle rates stayed stubbornly, irrationally high, barely moving while the cost of money moved underneath them by enormous margins. A 2022 review of decades of data described them as very sticky. Work in

2024 and beyond found them high still. So the 1987 result is not a relic at all. It is the first data point in an unbroken line running to the present.

That is what I am willing to call battle-tested. Not old, so trust it. Re-tested across four decades and two interest-rate regimes, so trust it.

The same discipline has to cut the other way, and this is the part I care about most. Where a finding is old and has *not* been convincingly re-tested, or worse, where later work has failed to reproduce it, I say so and I don't lean on it. There is a famous result about escalation of commitment, the tendency to throw good money after bad into a failing project, that fits this book's argument almost too neatly. It is from 1976, it is canonical, and recent attempts to reproduce it in the lab have struggled, with some finding the opposite effect. So when it turns up later it turns up with that warning attached, doing lighter work than its fame would allow. A durable pattern earns confidence. A famous but shaky one does not, however much I might want it to.

The distinction matters because the reliability of the whole book rests on it. Battle-tested is a badge I am going to hang on specific findings that earned it by re-test and refuse to hang on ones that didn't. Applied to everything old it would mean nothing. Applied only to what survived, I think it means quite a lot.

. . .

Why this makes the AI question sharper, not softer

Here is the payoff, and it sets up the hinge the whole book swings on.

If the pattern were new, a quirk of generative AI or a symptom of this particular technology, then the sensible response would be to wait it out. New things settle. Frontier tools mature,

governance catches up, the exception gets absorbed into the rule, and in five years the whole shadow-AI phenomenon looks like an adolescent phase the industry grew out of.

But it isn't new. It has survived every previous attempt by every previous technology to either enable it or stamp it out. It outlived the mainframe, the PC, client-server, the web and mobile. It shrugged off 3M's programme and Google's programme and every governance crackdown in between.

Which turns the question you actually care about into something much more pointed. Not "will this shadow-AI thing blow over," because the historical record says it won't, any more than the previous five waves blew over. The real question is the one this entire book is walking toward:

Is AI just the sixth wave of a pattern that has survived all five before it, or is it, for the first time, the thing that changes the pattern itself?

That question is only available to us because the research is old. A book built on 2025 citations couldn't ask it, because it would have no way to establish that there was a sixty-year pattern for AI to either continue or break. The very thing that looked like this argument's weakness, the dates, is what lets it ask its most important question with a straight face.

We will spend the book earning an answer. But we cannot even ask it honestly until we have done the next thing, which is to understand why the machine buries good small ideas in the first place, reasonably and defensibly and catastrophically, at every single step.

That is Part II. It starts with the gate.

. . .

Surface Check

So — *should Maria raise her hand on Tuesday?*

Still not yet, and now we know she's in old company. The champion in 1963, the engineer with the unapproved spreadsheet in 1995, the employee marked down for a shadow tool that worked — Maria is the current name for a very old boat. That tells us the fear is rational and the pattern is permanent. It still doesn't tell us what happens in the room. For that we have to see the machine itself. Ask again after Part II.

INSTRUMENT

The Durability Ledger

A short one, and it is a habit rather than a worksheet.

The next time somebody in your organisation describes a problem as new — the AI shadow-tool problem, the remote-work-visibility problem, the citizen-developer problem — try writing it on the left of a page, and on the right, the oldest version of the same problem you can find. Not the same technology. The same shape: someone routing around the official way to get the work done, and hiding it.

You will almost always find one, and it will almost always be older than you expected. That gap, between how new the problem feels and how old its shape actually is, is in my view the single most useful diagnostic in this book. A problem that is genuinely new might be solved by a new tool or a new policy. A new *name* for a sixty-year-old problem will not be, and the programme currently promising to solve it is, with high probability, the latest in a long line of programmes that didn't.

Keep the ledger for a month. It will change how you read every digital transformation deck you are ever handed again.

Next: Chapter 3 — why a pattern this old is arriving now at a scale it has never had, which turns out to be two changes rather than one.

CHAPTER THREE

Why Now

Five years before Maria built anything, she sat in a room and watched a man from a consultancy explain her own plant to her.

He was good, and that was the uncomfortable part. He had the diagrams, the process maps, the failure taxonomies. He knew which supplier codes mapped to which line, and he knew it because he had spent four months building the model that mapped them. He was thirty-one and he had been in the building since March.

Maria had been there nine years.

She knew things he didn't. She knew the seal supplier's quality dipped every August, that the third shift logged deviations differently to the first, that the downtime on line two had been logged against the wrong press since a system migration nobody ever fully unpicked. None of that was in the deck, though, because the deck was the official version. And when the programme ended and the man went back to wherever men like that go, the deck stayed, and the model stayed, and the understanding of how the model actually worked left the building in his laptop bag.

Two years later somebody asked Maria why the downtime on line two looked strange.

She said she'd raise a ticket.

. . .

What it feels like to work in a hollowed-out company

I want to spend some time on that feeling, because it is the part the research handles worst and the part every engineer in a large company recognises instantly.

It is not incompetence. Maria is not incompetent, the plant is not badly run, and the consultancies were not con artists. Every individual decision in the sequence was defensible, which is a sentence you should get used to, because it is the shape of every mechanism in this book.

What it is, I think, is something more specific and more corrosive: being accountable for outcomes you no longer have the means to produce.

You are the quality engineer. Quality is your name on the org chart. But the system that measures quality was built by a vendor, is maintained under a support contract, and can only be changed by raising a request that goes into a queue managed by a team in another country who will ask you to justify the business impact of the change in a form with eleven fields. You know exactly what the change should be. You have known for a year. The knowledge in your head has no route into the system that runs your job.

So you develop the habits. You keep the spreadsheet. You learn which of the three official reports is the one that is actually right. You build a private mental map of where the data lies to you, and you carry that map around as tacit knowledge, and you never write it down, because writing it down would take a week and nobody asked for it and there is no field for it in the eleven-field form.

Without ever deciding to, you become a person who works around their own systems.

And here is what that does over nine years, which is the thing I want you to sit with before we get to the evidence. It teaches you that building things is not what people like you do. Building is what the vendor does. What people like you do is specify, request, escalate and wait. You become the customer of your own operation.

That belief is the ground this book is planted in. It is not a skills gap. It is a category belief about who gets to make things.

. . .

How the knowledge left

The path by which companies arrived here was not stupid either. At every step it was financially sensible.

Through the nineties and two-thousands, large organisations worked out that a great deal of what their internal IT departments did was undifferentiated. Running a mail server was not a competitive advantage. Maintaining a data centre was not a competitive advantage. And the people who did it were expensive, hard to hire and hard to retain.

So it went out. First infrastructure, then applications, then whole functions. The SaaS wave completed the logic, because why maintain a system at all when you can subscribe to one? Each decision reduced cost, transferred risk, and freed the organisation to focus on its actual business. Each one was correct on its own terms.

What went out with the work, though, was the learning.

This is the part that never shows up on the business case, because it is not a line item and it does not arrive for years. When a vendor runs your system, the vendor's engineers accumulate the experience of running it. They learn what

breaks, why it breaks, what the workarounds are, how the thing actually behaves under load on a Tuesday in November. That learning compounds, on their side of the contract.

Meanwhile your side accumulates a different competence: procurement, contract management, vendor governance. Real skills, and useful ones, and none of them the skill of building the thing.

Ten years of that and you have an organisation full of capable, senior, well-paid people who are extremely good at commissioning work and have not personally made anything in a decade.

Then somebody stands up at a town hall and says the company must move fast on AI.

. . .

What the research actually establishes

The academic version of this has been around since 1992, which by now you will recognise as the book's usual pattern.

Richard Bettis, Stephen Bradley and Gary Hamel published "Outsourcing and Industrial Decline" in the *Academy of Management Executive*. Their argument was that outsourcing decisions, individually rational, could cumulatively hollow out a firm's ability to compete, not through any single bad choice but through the steady transfer of capability to suppliers who then owned it.

The finding has been sharpened since. Masaaki Kotabe and Michael Mol studied manufacturing businesses and found the relationship is not a straight line at all but a curve with a top to it. Outsource nothing and you carry costs you did not need to carry and lock yourself out of things other people simply do better. Outsource more and performance improves, for a while. Past a certain point it turns and comes back down, and it comes

down for a different reason than it went up — not because any particular decision was wrong, but because the firm has cumulatively handed away things it turned out to need. Somewhere in the middle there is a peak, and the curve is worse on both sides of it.

Two properties of that shape matter more than the shape itself. The first is that the descent is built out of the same decisions as the climb. Each one resembles the last, each arrives with a business case that is honestly argued and probably correct in isolation, and none of them is the mistake. The mistake is the sum, and nobody's job is the sum. The second is that the peak is not marked. Nothing announces the crossing. Performance does not fall the following quarter, no committee is informed, and the plant that has just outsourced the last person who knew how to modify a line does not feel any different on the Monday morning. You discover where the top was by standing some way down the far side of it, years later, when somebody asks the organisation to do a thing it used to be able to do and the answer comes back that this is something we buy.

The most pointed result for our purposes is Carmen Weigelt's, in the *Strategic Management Journal*. Outsourcing buys a firm access to a new technology. It does not buy the ability to integrate that technology with the way the firm already works, or to build on it afterwards — and Weigelt's term for that ability, *integrative capability*, is internal by definition. Access can be purchased. Absorption has to be grown.

Her setting is banks adopting the internet, not anything with a factory in it, and I would rather say that than have you find it out later. What she found there was that heavier reliance on outsourcing carried a performance cost, and — this is the part worth carrying out of the paper — that the cost was smaller for firms that already had experience with a related technology. Prior internal capability did not become unnecessary once the

work went outside. It became the thing that decided whether going outside hurt.

Read that twice, because it inverts the intuition the whole outsourcing era ran on. You do not save the capability by buying the service. The capability is what determines whether buying the service works at all, and the conditions under which it is grown are precisely the ones you have just handed away.

Sean Handley documented what follows as the perilous effects of capability loss: firms that lose internal competence manage their outsourcing relationships worse, which produces worse outcomes, which typically produces more outsourcing. The loop closes.

Underneath all of it sits the resource-based view of the firm, which holds that competitive advantage comes from resources that are valuable, rare and hard to imitate. Capability is exactly such a resource, and it is not a thing you own like a building. It is a thing you maintain by using.

. . .

The town hall

The research I have just cited stops at the edge of the company. Firm capability, firm performance, competence measured in numbers that hold at the scale of the whole organisation. It never comes down the stairs, onto the floor, into the particular Tuesday when the plant manager stands up in front of two hundred people and says the words everyone has been waiting for, which are that we need to move fast on AI.

Watch Maria in that room.

She is not excited. She notices that she is not excited, and it bothers her, because five years ago she would have been. Around her some people are nodding the way you nod at a slide. A few

are typing. Nobody looks the way the plant manager plainly hopes they will look.

What Maria feels, sitting there, is closer to being asked to run a race after a decade of being driven everywhere. She has spent nine years becoming excellent at a particular thing, which is knowing this plant, its suppliers, its lies, its August dip, and almost none of that excellence is the kind you can act on directly anymore. When something needs building, someone is hired to build it. When something needs changing, she raises a request. She has become very good at describing what should happen and completely out of practice at making it happen. And now the manager wants speed, and boldness, and initiative, from a room full of people the company spent ten years teaching to wait.

She could tell him why the enthusiasm isn't there. She won't. But she knows.

. . . .

That much I will state without a footnote, because you can walk into any hollowed-out company and find it in the third row of the town hall: the capable, senior person who has quietly concluded that building things is not what people like them do anymore. It is not rare and it is not subtle once you know to look for it. In my experience it is the single most common expression in the room when the AI announcement lands.

What I cannot tell you, and what Maria herself couldn't tell you if you stopped her on the way out, is exactly why the excitement drained away. Whether it is the habit of specifying instead of making. The slow discouragement of a hundred tickets raised into a queue. An honest reading of skills that have gone quietly rusty. Or something nearer to identity, some sense of what her job is for that shifted so gradually she never caught

it moving. Probably all of them together, braided, pulling harder on some people than on others.

I have a guess, and I will give it as a guess. But the guess is not the point. The point is that the dread is real and observable and everywhere, and that it arrives at the exact moment the company most needs its opposite. You do not need to resolve the why to see that the what is going to be a problem.

Because here is what Maria does after the town hall. She goes back to her desk. She does not propose anything. She does not raise her hand about the tool she has already built, the one that works, the one that is exactly what the plant manager stood up and asked for.

She opens the browser tab, and she gets back to work.

. . .

The second thing that changed

Now the other half, and this one moves fast.

For most of the sixty years we traced in Chapter 2, running a submarine was expensive. Schön's champions in 1963 needed political capital and a laboratory. Lockheed's Skunk Works needed a building, a budget and an executive who would look away. The engineer with the unauthorised spreadsheet in 1995 needed weeks of evenings and a real command of the tooling.

The cost of hiding was high, and the cost of building while hidden was higher, which is why the fleet was historically small. It took a certain kind of person: stubborn, senior enough to survive being caught, technical enough to do it alone.

That constraint is gone.

The MIT numbers from Chapter 1 are the measurement of its absence, with around ninety per cent of employees using personal AI tools for work against forty per cent of companies with sanctioned ones. That ratio is not really a story about

enthusiasm for a new technology. It is a story about a barrier that fell.

Maria's tool cost six hours and eighteen euros a month. In 1995 the same capability — parse unstructured reports, classify failure modes, cluster by supplier, flag recurrences — would have been a project. A real one, with a budget, a spec and a vendor. It would have gone through the gate we examine in the next chapter, and it would have failed there, and we would never have heard about it.

Now it is a weekend.

Notice what specifically changed, because it is not quite what people say. AI did not make people want to build things they were not allowed to build. They always wanted that, and Chapter 2 is sixty years of them doing it. What AI changed is the economics of deviance. It took an activity that required real resources and made it available to anyone with a browser and an afternoon. The behaviour did not change. The price did.

There is a second property that is easy to miss, and the research on covert AI use names it directly. When Waters-Lynch and colleagues studied hidden generative-AI use inside firms, they identified concealability as a condition shaping whether hidden work happens at all, meaning how easily the activity can be kept out of sight. A skunkworks needs a building, and you cannot hide a building. Maria's workflow lives in a browser tab on a personal account and produces no artefact anyone else can see. It is not merely cheap to build, it is nearly free to conceal.

Cheap to build, cheap to hide. Both barriers fell at once.

. . .

The convergence

Put the two halves together, because separately they are interesting and together they are the reason this book exists.

The capability to build things independently arrived, cheap and fast and available to anyone, at precisely the historical moment when large organisations had spent thirty years training their people to believe that building things was somebody else's job.

Those two facts are in tension and the tension has to resolve somewhere. What it produces is not a workforce that suddenly starts proposing ambitious projects, because proposing is the part they have learned doesn't work. What it produces is a workforce that quietly solves its own problems and doesn't mention it, because the tools are now within reach of the individual while the institutional route remains exactly as closed as it ever was.

That is the fleet, and that is why it is so much larger than it has ever been, and why it grew so suddenly that most executives have not registered it yet.

There is a bleaker way to state it and I think it is the accurate one. The organisation hollowed out its ability to build. Its people are now rebuilding that ability privately, on their own accounts, at their own expense, in their own time, and hiding it, because the organisation's formal machinery would treat it as a policy violation rather than as the restoration of a capability the organisation itself gave away.

Maria is not going around her employer. She is doing, at eighteen euros a month, the thing her employer stopped being able to do.

. . .

The question this leaves open

I want to close Part I by planting something and then leaving it alone for two hundred pages.

We have established that the pattern is old, structural, and survives every technology. We have established that it just got radically cheaper, and that the ground was unusually well prepared for it. What we have not established, and what I do not think anyone can yet establish, is whether AI is simply the sixth wave of a very old thing or the first one that changes what the thing is.

The case for the sixth wave is Chapter 2 in its entirety. Mainframes, PCs, client-server, the web, mobile. Each arrived, each was hailed as transformative, each got absorbed, and the pattern of hidden bottom-up work continued underneath, unbothered. On that reading, shadow AI is what shadow IT was, at greater scale, and it will be governed and absorbed the same way.

The case for different this time is harder to dismiss than usual. Previous technologies changed what employees could use. This one changes what employees can make, alone, without permission, at a cost that no longer requires approval. That is not a difference of degree in the tooling. It is a difference in who is capable of production.

I have a view. I am keeping it until Chapter 13, because you should watch the machinery work before you decide whether anything can break it.

What comes next is that machinery. Part II is the anatomy of how a large organisation, acting reasonably at every step, buries the best small ideas it has. It begins where every idea has to go, and where a surprising number of them stop.

. . .

Where the evidence stops

This chapter makes a historical argument, and historical arguments are the easiest kind to make badly, so it is worth marking the joins.

The hollowing-out evidence is narrow where the claim is wide. Kotabe and Mol's curve comes from manufacturing businesses, Weigelt's result from banks adopting the internet, and Bettis, Bradley and Hamel's warning was an argument made in 1992 rather than a measurement taken since. I am reading them together as one story about capability leaving firms over three decades. No single study says that, and nobody has run the study that would.

I am claiming a sequence, and sequences are where hindsight does its worst work. The order I have given — outsource, lose the people who knew, discover the loss when something needs changing — is the order the sources support individually. It is not an order anybody has traced through one company over thirty years, and the reason it reads so cleanly is partly that I have arranged it. Every account of a slow decline looks inevitable once you know it happened.

The convergence is my argument, not a finding. That the hollowing-out and the arrival of cheap capability meet now, in this decade, at this cost point, is the load-bearing claim of the chapter and there is no study behind it at all. It is two well-evidenced trends and a line drawn between them by me. If it is wrong, the most likely way it is wrong is that the capability was never as thin as the outsourcing figures make it look, because the people who knew things stayed and simply stopped being asked.

And this is the chapter with the least direct evidence in Part I. Chapters One and Two rest on mechanisms with controlled work behind them. This one rests on economic history, which is

real but is not an experiment, and on a reading of the present that will look either obvious or foolish in ten years with nothing in between.

. . . .

Surface Check

So — should Maria raise her hand on Tuesday?

Still not yet, and the reason has moved. In Chapter 1 we did not know enough about her. Now we know exactly why she is carrying this alone: an organisation that spent thirty years learning to buy capability rather than keep it, and a person in it who absorbed the lesson that building is somebody else's job. She broke that lesson quietly, and it worked. What we still can't see is what the room does with an idea when it actually arrives — whether her fear is calibrated or just inherited. That's Part II, and it starts with the gate.

INSTRUMENT

The Cost-of-Deviance Curve

An exercise in three columns. Take one improvement in your area, a real one, something a single person could conceivably build, and estimate what it would have cost to create *and* to conceal in three eras.

1995, the shadow IT era. What would this have taken? Probably weeks of evenings, real technical depth, and something visible — a database, an application, a machine under a desk. Concealment cost: high. Someone would notice.

2015, the SaaS era. A cloud tool, a corporate card, maybe a departmental subscription nobody scrutinised. Cheaper to build,

and moderately concealable, until procurement or security ran an audit.

Today. A browser tab, a personal account, an afternoon. Concealment cost close to zero, because there is no artefact and no footprint on any system your organisation monitors.

Now the question the exercise is actually for. As the cost of building-and-hiding fell by orders of magnitude, did your organisation's process for surfacing an idea get any cheaper?

For almost everyone the answer is no. The approval path costs roughly what it cost in 1995, with the same forms, the same committees and the same waiting. Which means the gap between what it costs to do it quietly and what it costs to do it properly has widened enormously, and is still widening.

That gap is the fleet. It is not a culture problem. It is an arithmetic one, and your organisation is on the wrong side of it.

Next: Part II. The machinery — beginning with the gate.

PART II

Why the System Buries the Best Ideas

Everything so far has been about the fleet: that it exists, that it is old, that it just got cheap, and that the people crewing it have good reasons to stay down.

Now we go to the surface and look at what is waiting for them there.

For the next three chapters I am going to take the machine apart, meaning the actual mechanism by which a large, well-run organisation buries the best small ideas it has. I want to be clear at the outset about the spirit of it, because it would be easy to read Part II as an exposé and it is not one. There are no villains in here. Every part of the machine I am about to describe was built by intelligent people for defensible reasons, and every part of it, taken alone, does something useful. That is precisely the problem. A machine made entirely of reasonable parts can still produce an unreasonable result, and this one does, reliably, in companies that have done nothing obviously wrong.

We start where every surfaced idea has to start. Maria, in Chapter 3, chose not to raise her hand. This chapter is about what the room does to the ideas of the people who do.

CHAPTER FOUR

The Gate

Two ways to die

Imagine two people walk into the same review meeting on the same morning.

The first has a genuinely new idea, a way of doing something the company has never done, using a method nobody in the room recognises. Call her the radical.

The second has a small improvement to something that already exists, a better way to run a process the company runs every day, that saves a bit of time and money and steps on nobody's strategy. Call her the incrementalist. This is Maria.

Both of them will very likely walk out with a no. And here is the thing that took me an embarrassingly long time to understand, and that most of the writing on corporate innovation still gets wrong: they will get their no for completely different reasons. The radical and the incrementalist are killed by two different mechanisms, and confusing them is why so much well-meaning advice about fostering innovation misses.

Almost everything written about innovation being stifled is about the radical, and about the bias against the genuinely new. That bias is real, and we will spend the first half of this chapter on it, because it is well documented and it matters.

But it has nothing to do with why Maria's idea dies. Maria's idea is not novel, and that is the whole point of it, because it is a sensible, legible, low-risk improvement to a known process. The anti-novelty machinery does not even engage. Something else kills her idea, something the innovation literature almost never names, and it is the more important half of this chapter because it is the half that explains the fleet.

Let us take them in order.

. . .

The first death: the bias against the new

Start with the radical, because hers is the famous story.

There is a comfortable belief, especially among leaders, that organisations want creativity — that the new idea's problem is a shortage of it, and that if people would just be braver and think bigger, innovation would follow. It is a nice belief and it is almost the reverse of what the evidence shows.

In 2012 Jennifer Mueller and colleagues ran a study with a title that gives the game away, “The Bias Against Creativity.” What they found is that people hold an implicit bias against creative ideas, a genuine and measurable aversion, that surfaces precisely when they are feeling uncertain. And a decision meeting evaluating a novel idea is, by definition, a room full of uncertainty. The more genuinely new the idea, the more uncertainty it generates, and the more the bias fires. So the ideas that are most creative are the ones most reliably rejected, by the same people who will tell you sincerely, over coffee afterwards, that they wish their teams were more creative. They are not lying. They simply cannot feel the bias operating, because it runs underneath the reasons they give themselves.

The finding has held up and sharpened since. Kevin Boudreau, Eva Guinan, Karim Lakhani and Christoph Riedl ran a grant

proposal process at a research university and randomised which evaluators saw which proposals, generating over two thousand evaluator–proposal pairs. Two things came back, and both are quietly brutal. Evaluators marked down proposals that were highly novel. They also marked down proposals close to their own expertise. Expertise did not correct for the novelty penalty; it carried a penalty of its own. Novelty was not a tiebreaker in the idea's favour. It was a penalty.

And Paola Criscuolo, Linus Dahlander, Thorsten Grohsjean and Ammon Salter — Criscuolo and Salter again, from the bootlegging study in Chapter 1, this time with two new collaborators — took the question inside the corporation, studying how the panels at a large professional services firm actually chose which R&D projects to back.

What they found was not a simple prejudice against the new, and the shape of it matters. Some novelty helped: a proposal that departed moderately from what the firm already did received a larger share of the funds it asked for than a wholly conventional one. Past a point the curve turns over, and too much novelty was selected against, because it was difficult to appreciate. The ideas that lost out were not the mediocre ones. They were the ones furthest from what the panel already knew how to think about.

And where the curve turns depends on the panel rather than on the proposal. A panel carrying a heavy workload became less willing to back novelty. A panel with a spread of expertise became more willing. The same proposal, in other words, is worth different amounts depending on who is in the room and how much else they have to get through that day.

So the gate is not a neutral filter that occasionally lets a good radical idea through. It has a thumb on the scale, the thumb presses hardest on exactly the ideas that depart most from what

is already understood, and how hard it presses is a property of the room rather than of the idea.

Why? Not stupidity and not cowardice, which is the usual disappointing answer. A genuinely novel idea is by definition the one with the least evidence behind it, because nobody has done it before. It is the hardest to justify, the hardest to forecast, the easiest to attack, and the one most likely to make the person who championed it look foolish if it fails. An evaluator behaving with perfect rationality about their own risk will mark it down. The bias is not a malfunction of the evaluation process. It is the evaluation process, working as designed, applied to something the design cannot price.

That is the radical's death, and it is worth understanding. But it is not Maria's death, and if you only understand this one you will misdiagnose the fleet completely. You will conclude that the problem is fear of the new and you will prescribe more courage, more vision, more tolerance for moonshots, none of which would have saved Maria's supplier-deviation tool. Her tool is not a moonshot. It is the opposite of one. It is a small, ordinary, sensible improvement to a process the plant already runs every day, the kind of idea nobody would call brave or visionary or creative. And that ordinariness, which ought to be exactly what makes it safe to surface, turns out to be irrelevant to whether it survives, because a different mechanism is waiting for it.

. . .

The second death: territory

Here is the mechanism nobody talks about, and once you see it you cannot unsee it in any large organisation you ever walk into again.

Maria's idea is small, sensible and low-risk. It threatens no strategy and generates almost no uncertainty. By every measure in the anti-novelty literature it should sail through.

And it dies anyway. Ask yourself why, honestly, and the answer is not about the idea at all. It is about whose idea it is, and whose patch it lands on.

Because Maria's tool reviews supplier deviations, and there is a team, call it Quality Systems, whose job is supplier deviations. That is their remit. Their budget line. The thing they are staffed to own. And Maria, a quality engineer three rungs down and one department over, has just built something that works inside their territory, without asking, and it is better than what they have.

Watch what happens when she surfaces it. The Quality Systems lead does not say this is a bad idea. He can't, because it is plainly a good one, and saying otherwise would expose him. What he says is subtler and you have heard it a hundred times. *Interesting, though this really sits within our roadmap. Or: We looked at this, actually, and it's more complex than it appears at the plant level. Or the purest form, the one that ends things without a single word of disagreement: Great initiative. Let's make sure it's aligned with the programme.*

None of those is a rejection of the idea. Every one of them is a defence of the border.

What the research calls this

For a long time this had no proper name in the management literature, which is part of why it goes undiagnosed. It was folk knowledge. Everyone who has worked in a big company knows about turf, but folk knowledge is easy to wave away as office politics, a personality problem, a few insecure managers rather than a structural force.

Then, in 2005, Graham Brown, Thomas Lawrence and Sandra Robinson published a paper in the *Academy of Management Review* that gave it a spine. They called it territoriality, and their central move was to show that people in organisations become territorial not only over physical space, the office and the parking spot, but over ideas, roles and responsibilities. The intangible territory. The remit. And they mapped the behaviours that defend it: marking, where you signal that something is yours, and defending, which comes in two flavours, anticipatory, heading off an incursion before it happens, and reactionary, repelling one that already has.

Maria's surfaced tool is an infringement in the exact technical sense, and what she triggers in the Quality Systems lead is reactionary defending, a response that the research shows has essentially nothing to do with the merit of what infringed. Brown and Robinson followed this up in 2011 in *Organization Science*, looking directly at how people react to territorial infringement, and the reaction is what you would expect from the word. The interloper is repelled and the border reasserted, regardless of whether the interloper was carrying something valuable.

But the finding that should stop you is a later one. In 2015 Graham Brown and Markus Baer published work in the *Journal of Applied Psychology* with a title that could serve as a subtitle for this whole chapter, "Protecting the Turf: The Effect of Territorial Marking on Others' Creativity." They showed experimentally that when people mark their territory it measurably suppresses the creativity of everyone else around it. The mere act of one party staking a claim causes other people to hold back, contribute less, stop bringing things forward.

Read that against the fleet and it lands hard. The Quality Systems team does not have to reject a single one of Maria's ideas for the territorial effect to work. Their ownership of the

patch, marked and visible and understood by everyone, is already, before Maria builds anything, a reason for her to keep it quiet. The border does not just repel the ideas that reach it. It stops most of them from ever being surfaced at all, and territoriality therefore does its heaviest work invisibly, in the ideas that die in the mind of someone who correctly reads whose patch they would be stepping on.

The old, honest version of it

None of this is new, of course, since nothing in this book is. The territorial reflex has a well-worn cousin you will already know by name, which is Not Invented Here, the reflexive rejection of ideas that came from outside the group.

Ralph Katz and Thomas Allen documented it properly back in 1982, studying fifty R&D project groups. They found that stable, long-tenured teams developed the belief that they possessed a kind of monopoly of knowledge in their area, and used that belief to reject outside ideas, to the measurable detriment of their own performance. They were hurting their own results to defend their ownership of the problem, and they could not see themselves doing it.

Later work by Dirk Antons and Frank Piller framed it precisely: Not Invented Here is best understood as an attitude-based bias, a psychological mechanism that leads people to reject external knowledge in an irrational, economically suboptimal way, regardless of how good the knowledge actually is. The value on offer does not enter the calculation. The only thing that matters is that it came from outside.

Which is Maria exactly. Her tool does not fail a cost-benefit test, it never reaches the cost-benefit test. It is disqualified upstream of merit by the single fact that it originated in the wrong place: three rungs down, one department over, outside the walls of the team that owns the problem.

Why this is the more important death

I have spent longer on the second mechanism than on the first, and that is deliberate, because it corrects what I think is the single most common mistake in how organisations understand their own buried innovation.

The anti-novelty bias is real, it is what everyone points to, and it kills radical ideas. But radical ideas are rare. Most of the value trapped in the fleet is not radical at all. It is Maria's kind, meaning sensible, incremental, low-novelty, high-return improvements to things the company already does. Those ideas are not killed by the bias against creativity, because they are not creative in the sense that triggers it. They are killed by territory.

And this matters enormously for what you do about it, which is the whole back half of this book. If you believe the problem is fear of the new, you will invest in courage — innovation programmes, moonshot funds, permission to fail — and you will watch the fleet stay submerged, because you will have treated a disease the fleet mostly does not have. The increments are not hiding because the company is afraid of bold ideas. They are hiding because the company is organised into territories, and every good idea that crosses a border reads, to the border's owner, as an incursion to be repelled, no matter how much value it happens to be carrying.

There is a second consequence and we will pick it up two chapters from now, but let me plant it here. Territory explains why the increment specifically submerges, and it does so without needing the idea to be novel. That closes a gap the standard story cannot close. If the only mechanism were fear of novelty, low-novelty increments would surface freely, and they don't. Territory is why.

The gate is not the meeting

One last thing, because it is the piece that makes everything after this chapter make sense.

When I say “the gate,” it is tempting to picture the review meeting, the stage-gate, the steering committee, the formal go/no-go. Those are real and they do their share of the killing. But the deeper point of this chapter is that the gate is mostly not a meeting at all. The gate is the anticipation of the meeting.

Maria never presented her tool. There was no review, no panel, no formal no. The gate did its work anyway, perfectly and silently, because Maria could read the territory from her desk. She knew whose patch it was. She knew what “let’s align it with the programme” means. She had watched it happen to other people. The bias against novelty and the defence of territory did not need to be applied to her idea in a room, because she had already applied them to herself, in advance, and drawn the correct conclusion.

That is what a gate is, fully understood. Not a checkpoint that rejects ideas, but a pressure that stops them being offered at all. The rejections you can see are the small visible tip of it. The real gate is the one running in the heads of everyone who never bothered to walk up to it, because they already knew what it would say.

Which raises the question the whole back half of the book has to answer, and I will leave it hanging here. If the gate mostly operates by anticipation, then you cannot fix it by reforming the meeting. You have to change what people anticipate. Hold that thought, because it gets a good deal harder before it gets easier — next we leave the gate and go to the place the anticipation

comes from, which in a manufacturing company is a very specific and very expensive kind of decision.

. . .

Surface Check

So — should Maria raise her hand on Tuesday?

Now we know what the room would do, and it isn't what the innovation books would predict. Her idea wouldn't die because it's too bold — it isn't bold at all. It's a small, ordinary improvement, the safest kind of thing to bring forward. It would die because it lands on someone else's patch, and the someone else would defend the border without ever having to call the idea bad. Her instinct to stay quiet wasn't cowardice. It was an accurate reading of the territory. So the fear is starting to look calibrated — but we still don't know why the whole organisation is tilted to protect the big and the owned over the small and the useful. That tilt has a source, and in a factory it's made of concrete and steel. Next.

INSTRUMENT

The Gate Audit

Not a process map. A set of questions you answer honestly about the last good small idea that reached you, or that you brought to someone else.

1. **Did it die at a meeting, or before one?** If before, you are looking at anticipation rather than evaluation, and no amount of fixing the meeting will help. Most deaths are before.
2. **Was it rejected for what it was, or for where it came from?** Be honest here. “It needs to align with the roadmap” is a

where-it-came-from rejection wearing a what-it-was costume.

3. **Whose patch did it land on, and did that person get a say?** If the idea touched a team's remit and that team was in the room, the merit of the idea was probably not the deciding variable.
4. **Was it novel, or just new to them?** These are different deaths. A genuinely novel idea hit the bias against creativity. A sensible improvement that got killed hit territory. The fixes are not the same.
5. **How many good small ideas do you think you never hear about at all?** You cannot count these directly, which is rather the point. But the ratio of ideas killed in meetings to ideas never surfaced is, I would argue, the single most important number in your organisation, and it is almost certainly lopsided in a direction that should frighten you.

The audit's real purpose is question five. Everything above it exists to make you believe the invisible number is large. It is.

Next: Chapter 5. The Big Factory — where the tilt comes from, why the doors only swing one way, and the slow erosion nobody watches.

CHAPTER FIVE

The Big Factory

There is a mountain in the Alps that Maria's company put on a brochure once, a clean triangular peak, the kind a child draws. The brochure was making a point about permanence. *We have been here for over a century*, it said, more or less. *Solid as the mountain.*

The comparison is better than the brochure knew, though not in the way it intended.

Because the mountain is not permanent, and it is not solid, and it is not particularly a monument to the dramatic event that made it. The peak everyone photographs was thrust up in a geological instant, a violent spasm of tectonic collision, the kind of event that gets into the textbooks. But that thrust is not what gives the mountain its shape today. What shapes it now, continuously, is the opposite of an event. It is erosion: rain, ice, wind and gravity, removing a few grains at a time, every second, for millions of years. The dramatic uplift built the raw mass. The invisible grinding carved everything you would recognise.

Nobody watches the erosion. You cannot watch it, because it moves too slowly to see. If you stood on that peak for your entire life and stared, you would perceive nothing changing at all. And yet the erosion is winning, and it always wins. Every

mountain range has a life cycle that ends the same way, with the uplift brief and finite, the grinding patient and relentless, and given enough time the grinding takes the whole thing back down to a plain.

A large manufacturing company is a mountain in exactly this sense. Its shape is made by two forces on wildly different timescales, and it systematically mistakes which one matters.

. . .

The uplift

Let me start with the force everyone can see, because it is the one the whole organisation is built to serve.

A car plant is a staggering object. Not just large, but committed. When a manufacturer decides to build or retool a factory it is spending a sum with a lot of zeroes, on equipment that will be bolted to a specific floor in a specific country for a specific purpose, and that cannot meaningfully be undone. You do not half-build a body shop. You do not retool a paint line as an experiment and roll it back next quarter if it disappoints. The press that stamps door panels weighs as much as a house and was designed around the exact geometry of a car that will be in production for years.

These are the uplift events. Enormous, rare, dramatic, and irreversible, which is the word that will run through the rest of the chapter.

The economists have a precise language for why irreversibility changes everything, and it is worth having, because it explains behaviour that otherwise looks like simple cowardice. In the late 1980s and early 1990s, Avinash Dixit and Robert Pindyck worked out the formal consequences of investing under uncertainty when you cannot take the investment back. Their key insight is deceptively simple: when an investment is irreversible,

committing to it kills your option to wait, and that option has value. The ability to hold off, to see how things develop, to not be locked in — that flexibility is worth something real, and an irreversible commitment spends it. So a rational manager facing a one-way door does not just weigh the cost against the return. They add a third term, the value of the option they are about to destroy by walking through it, which means they demand a much higher return before committing and they wait a good deal longer than a naive calculation says they should.

You can measure this waiting. Firms use hurdle rates, the minimum return a project must clear to get funded, and the striking, well-replicated finding is that these rates sit far above what the financial theory says they should. Companies routinely demand returns three or four times their actual cost of capital before they will commit, and as we saw in Chapter 2, in one of the sturdiest findings in the whole book, those rates barely move even when interest rates collapse. The bar is high, it stays high, and it exists precisely because the doors only swing one way.

There is a second reason the big commitment is the only kind that feels natural, and it is physical rather than financial. Investment in a factory is lumpy. It does not flow in smoothly, it arrives in enormous discrete lumps. The plant-level data on this is unambiguous. A manufacturing plant will go years changing very little and then, in a single concentrated burst, replace a huge fraction of its capital at once. Doms and Dunne, studying thousands of US plants, found that around half of a plant's investment over a fifteen-year period was compressed into just two or three consecutive years, with the rest of the time spent in relative stillness. The factory does not evolve continuously. It sits, and then it heaves.

Put irreversibility and lumpiness together and you get an organisation whose entire decision-making apparatus is shaped around the big, rare, one-way bet. That is what its senior people

are trained to evaluate. That is what the committees, the hurdle rates, the capital-approval processes and the whole expensive machinery of governance exist to handle. And it is right to exist. I want to be very clear about this, because the easy version of this book would cast the big-bet machinery as the villain and it isn't. When the decision genuinely is a house-sized press bolted to a floor for a decade, you *should* deliberate for months, demand a high return, and guard the door. That caution is correct.

The problem is not the caution. The problem is where the caution gets applied.

. . .

The one-way door and the two-way door

Here is a distinction that will do a lot of work for the rest of the book, and it comes, as it happens, from the world Maria's company is now being told to imitate. Jeff Bezos set it out in Amazon's 2015 letter to shareholders, and it has been quoted so often since that a good many of the people using it no longer know whose it was.

The idea is simple enough. Some decisions are one-way doors, hard or impossible to reverse, so you had better be sure before you walk through. Others are two-way doors, and if it doesn't work you walk back through and the cost of having tried is small. The whole art, the framing says, is telling them apart, because the two kinds of door deserve completely different treatment. One-way doors want deliberation. Two-way doors want speed, because the cost of hesitating exceeds the cost of a reversible mistake.

Now hold Maria's tool up against that distinction. Six hours to build, eighteen euros a month to run, on her own account. If it had failed she would have closed the browser tab and lost an

afternoon. It is a two-way door, and close to the purest two-way door imaginable: nearly free to try, nearly free to reverse, nearly free to walk back out of.

And here is the failure, stated as plainly as I can. The organisation only owns one kind of door-handling, and it is the wrong kind for Maria.

The manufacturer's every instinct, every process, every trained reflex is built for one-way doors, because that is what factories are made of. So when a two-way-door idea like Maria's approaches the machinery, it gets one-way-door treatment. It gets the deliberation, the demand for a business case, the request to prove the return, the alignment review and the wait, which is all apparatus that is correct for a house-sized press and catastrophic for a six-hour tool. You cannot ask a two-way door to clear a one-way-door hurdle rate. The tool that would have paid for itself in a fortnight is asked to justify itself as though it were a capital project, and it can't, and it dies. Not because the caution is wrong, but because the caution was built for a different kind of decision entirely and nobody ever separated the two.

This is the tilt Chapter 4 kept pointing at. The organisation protects the big and the owned, not out of malice or even preference, but because its entire decision-making muscle was built around irreversible bets and it applies that muscle to everything, including the reversible things that should never have been near it.

. . .

The resource nobody budgets

Capital is the tilt everyone can at least see. There is a second one, subtler and in some ways more decisive, and it explains something Chapter 1 left hanging.

Remember Maria's rival, the Quality Systems programme with eighteen months, two consultancies, a budget quoted in disbelief, and a roadmap slide that was, everyone agreed, very good. Nothing shipped to her plant. Maria's tool worked. And yet in any contest between them the programme would win effortlessly, without the two ever being formally compared. Why?

The answer is a resource that never appears on a budget line but is scarcer than money, which is attention.

Chapter 1 set out the mechanics: Ocasio's finding that a firm's behaviour is governed by how it allocates the attention of its decision-makers, and Bouquet and Birkinshaw's two routes to being noticed – weight, which is structural, and voice, which is effortful and can be spent. What matters here is the part that becomes specific to a factory.

Weight in a manufacturer is not a metaphor. It is the same lumpiness this chapter has just described, expressed in a room rather than in a capital plan. The Quality Systems programme carries budget, consultancies, an executive sponsor and a place on the transformation roadmap, and every one of those is a claim on the capital process that the previous section showed is built to serve exactly that shape. The mass that gets it through the gate is the mass that holds the room. It is one property doing both jobs.

Which is why the two were never in a fair fight, and never in a fight at all. A fight would require both to be in the room, weighed against each other. This is the sharper and crueller version of what Chapter 4 called territory: the increment does not lose the argument, because there is no argument. It is not defeated, it is unseen, and you cannot appeal a decision that was never made.

That is why the roadmap slide can be excellent for three years while nothing ships. The slide is not there to produce an

outcome. It is there to hold the weight, to keep the programme's mass parked in the centre of the room's attention, and it succeeds at that completely, which is the only thing it was ever measured on.

. . .

The clock

There is a third tilt, smaller than the other two but worth naming because it compounds them.

Manufacturers run on cycles. Budget cycles, model-year cycles, planning cycles, the great slow rhythm of a product that takes years to design and years to build and years to sell. This is not a flaw, it is a necessity given how long the physical bets take to pay back. You cannot run a car programme on a quarterly whim.

But innovation does not keep that time. A good idea does not arrive in Q3 aligned to the planning calendar. It arrives when it arrives, half-formed, needing a few weeks of quiet iteration to become anything, which is the delayed-assessment window from Chapter 1. And the cyclical organisation has nowhere to put a thing like that. It has a slot for next year's capital plan and a slot for this cycle's approved projects and no slot at all for a small useful thing that showed up in the wrong month and needs six weeks to prove itself. So the idea waits for a cycle that has no room for it, and by the time a cycle comes round that might, the moment has passed, and often so has the person's willingness to raise it again.

Capital tilts toward the big, attention tilts toward the heavy, and the clock tilts toward the scheduled. Three separate forces, three separate mechanisms, all pointing away from the small, the light, the unscheduled and the reversible, which is to say

away from exactly the kind of idea most likely to be trapped in the fleet.

. . .

The convergence, and a caution

I want to step back and name what has just happened, because it is the spine of the book assembling itself, and then I want to undercut it slightly, because that is the deal we made.

Across the last two chapters three scarce resources have each turned out to be biased in the same direction. Capital, because irreversibility and lumpiness make the big rare bet the only thing the machinery is built to handle. Attention, because weight pulls it toward mass and it is zero-sum. And careers and status, from Chapter 4 and returning in force in Part III, because the people who get promoted to allocate the first two are, as we will see, disproportionately the ones drawn to the big visible bet.

Three different currencies, and three different bodies of research that as far as I can find have never been introduced to one another. The investment-irreversibility economists do not cite the attention-allocation strategists, who do not cite the leadership-selection psychologists. And all three, independently, bend the same way. That convergence — the claim that these separate forces compose into a single systematic tilt against incremental innovation — is the argument of this book and it is mine. The individual findings are documented and I have cited them. The synthesis is the thing I am adding, and you should hold it as a strong argument rather than a proven fact, because no study proves it. It is an assembly.

And now the caution, which is real and which I would be dishonest to skip.

The whole capital story rests on the premise that manufacturing investment is enormous, lumpy and irreversible.

That has been overwhelmingly true for a century. But it may be becoming less true at the edges, right now. Flexible automation, additive manufacturing, modular and software-defined production lines are slowly making some manufacturing capital less monolithic and less irreversible than it was. A line you can reconfigure in software is not quite the one-way door a welded body shop is. If that trend runs far enough, the capital tilt I have just described could soften.

I do not think it undoes the argument, for two reasons. First, the softening is partial and uneven, because the house-sized press is not going anywhere soon and most of manufacturing is still built around genuinely irreversible bets. Second, and more importantly, capital is only one of the three tilts. Even if flexible manufacturing dissolved the capital story entirely, attention would still be zero-sum and would still flow to weight, and careers would still reward the big bet. The mountain would still erode in a direction the organisation cannot see. But I flag it honestly, because the strongest single leg of the capital argument is the one with a genuine question mark hanging over its future, and a reader who works in the most advanced and most flexible corners of manufacturing may already feel the floor shifting under it.

Which, if you want the optimistic reading, is rather the point of the whole book. The forces that bury the increment are not laws of nature. They are features of how organisations happen to be built, and things that are built can be built differently.

. . .

Back to erosion

Return to the mountain, because now the metaphor can carry its full weight.

The manufacturer is built, entirely and rationally, to serve the uplift, meaning the rare, enormous, irreversible bet. That is where its capital goes, where its attention pools, where its calendar bends, where its heroes are made. And the uplift matters, because it built the mass. But the uplift is not what determines the company's shape over time, any more than the tectonic thrust determines the mountain's. What determines the shape in the end is the erosion: the continuous, invisible, grain-by-grain accumulation of small changes. The incremental improvements. The learning. The thousand tiny adaptations that compound, unwatched, into the difference between a company that is quietly getting better every day and one that is quietly getting worse.

And the organisation cannot see its own erosion. It has no instrument for it. Every instrument it owns — every committee, hurdle rate, roadmap and budget cycle — is calibrated for uplift, for the discrete visible event. The erosion happens below the resolution of everything the company uses to perceive itself, which means the single most decisive force acting on a manufacturer over time is the one force its entire apparatus is structurally blind to.

That blindness has a name we have been circling for four chapters. It is the fleet. The erosion is the submerged innovation, the small compounding improvements happening constantly, below the waterline, in the accounts and afternoons of people like Maria, entirely unseen by an organisation whose every sense is tuned to watch for mountains being thrust into the sky.

We have now seen the machinery, and why it is built the way it is, and why every part of it is individually reasonable. What we have not asked is the question that should be nagging by now, the one a sceptic would have raised chapters ago. If all of this is

so well understood, if the erosion is so valuable, if the fixes seem so obvious, then why in sixty years has nobody fixed it?

That question turns out to have a darker answer than the machinery does. It is the whole of Part III.

. . .

Surface Check

So – should Maria raise her hand on Tuesday?

Now we can see the ground her fear stands on, and it's bedrock. Her tool is a two-way door in an organisation that only knows how to handle one-way doors, so it would be met with a caution built for something a thousand times its size. And even that overstates its chances, because to be over-cautioned it would first have to be seen – and it has no weight to pull the room's attention, while the programme that will never ship has all of it. Her idea wouldn't lose. It would go unnoticed. Raising her hand doesn't look brave from here; it looks like offering a grain of sand to a machine that only measures mountains. And yet the grains are what shape the mountain. Hold that contradiction – Part III is about why nobody has resolved it.

INSTRUMENT

The One-Way / Two-Way Door Sorter

Take the last ten ideas, requests or proposals that crossed your desk. For each one, ask a single question before any other: is this a one-way door or a two-way door?

A one-way door is expensive to reverse, or impossible. A hire. A platform migration. A press bolted to a floor. Get these wrong and you live with it.

A two-way door is cheap to try and cheap to undo. Most software. Most process tweaks. Most tools one person could build in an afternoon. Get these wrong and you close the tab.

Now the diagnostic question, and answer it honestly. How many of the two-way doors did your process treat like one-way doors? How many reversible, low-cost experiments were asked for a business case, sent for alignment, made to clear a hurdle rate, told to wait for the cycle?

Every one that was is a grain of erosion your organisation demanded be justified as though it were a mountain. And every person who watched that happen learned the lesson Maria learned, which is not to bother bringing the small useful thing, because it is not worth what it costs to get it through a door built for something else.

The number you want to drive toward zero is not the rejections. It is the mismatches, the two-way doors handled as one-way doors. That number is a direct measure of how much of your own erosion you are forcing to happen underwater.

Next: Chapter 6 — the increment that would pay for itself, and the arithmetic that means nobody will ever find out.

The Increment That Pays

In month six, Maria tried to fill in the form.

Not to submit it. That part matters. It was a Tuesday evening, the plant was quiet, and she had the internal template open in one window because she had been idly curious whether the thing she had built could be described in the language the company uses to describe things it pays for. A private experiment. Nobody was going to read it.

The template is not unreasonable. Every large company has one and they are more alike than their owners imagine, and this is one of the better ones. Two pages. Problem statement. Proposed solution. Scope and systems affected. Resource requirement. And then, about halfway down, the field that stopped her:

Expected benefit (quantified).

She could do part of it. The tool saves her roughly a day a week, which she could defend, because she had been counting: the deviation reports used to eat Thursday afternoon and most of Friday, and now they eat an hour. Call it eight hours, call it twenty per cent of an engineer, and there is a number a controller would accept. She typed something like that and looked at it.

Then she thought about the seal defect, and stopped.

Because the seal defect is the whole point, and the seal defect is not in that number. In month four the tool flagged a failure mode appearing three times in a fortnight across two different suppliers, which the old process would not have connected for another five weeks, because the old process was her reading reports in supplier order on a Friday evening and the two suppliers were nowhere near each other alphabetically. Five weeks of a defect not reaching the line. Anyone can price that. It is the single most valuable thing the tool has ever done and it is worth more than the day a week by a margin that is almost embarrassing.

And she could not have written it in the box. Not because it is hard to quantify after the fact. Because in month zero, sitting at her kitchen table pasting an anonymised report into a chatbot, she was not trying to catch seal defects. She was trying to survive Friday. Catching the seal defect was not a benefit she anticipated and then delivered. It was something the tool turned out to do, and she found out it could do it by having it, and running it, for four months.

There is a further problem, and it is the one she thinks about more.

The version that caught the defect is not the version she built. Her first version clustered by supplier, because that is how the reports arrive, how the system sorts them, how everyone thinks about them. It was fine. It saved her an afternoon. Some time around week six a batch came out wrong — she had pasted in a set from the wrong fortnight — and in the mess she noticed that the same failure mode was showing up under two supplier names, and that she had never once looked at the data that way in fourteen years, because nothing she had ever used presented it that way. So she rebuilt it to cluster by failure mode across

suppliers. That is the version that works. She got there by accident, in the sixth week, from a mistake.

You cannot put that in the box either. *Expected benefit* (*quantified*) wants a number at the point of asking, and everything that makes the tool worth anything was discovered after the point of asking, by running it.

Maria closed the template without saving. What she felt, sitting there, was slightly stupid – as though there were an obvious way to fill in the form and she was simply not clever enough to find it.

There isn't. The form was asking her, in month zero, for something that did not exist until month four. That is not a failing of hers and it is not a failing of the form. It is the shape of the whole problem, and this chapter is about it.

. . .

The oldest number in this book

Let me put the other half of the argument on the table first, because so far Part II has only shown you the machinery that buries things. It has not shown you that the buried things are worth anything.

They are, and the evidence is nearly ninety years old.

In 1936, T. P. Wright published an analysis of costs in aircraft manufacturing and reported a regularity that has been reproduced ever since. As cumulative output of an airframe doubled, the labour hours required to build one fell by something in the region of twenty per cent. Not because anyone invested in new equipment. Not because a programme was launched. Simply because the plant had built more of them, and the people in it had got better, in a thousand unremarkable ways nobody wrote down.

Kenneth Arrow gave it theory in 1962, in a paper about learning by doing. His argument was that a substantial part of productivity growth comes from experience accumulated in the act of production itself, and that this makes experience an input in its own right — one you cannot buy, order, or accelerate by spending, because the only way to accumulate it is to run the process and pay attention while you do.

The learning curve is now among the most durable findings in operations. It has been observed across aircraft, shipbuilding, semiconductors, chemicals, solar cells and surgical procedures. By the standard we set in Chapter 2 — not old, but *re-tested* — it is about as battle-tested as anything in this book. Every large manufacturer in the world has a version of it embedded somewhere in its planning assumptions, and everybody senior has heard of it.

The clearest statement of it I know did not come from an economist.

In 2017 Andy Jassy, then running AWS and now Amazon's chief executive, was asked by CNBC how he thought about competitors copying what his business had built. He said there is no compression algorithm for experience. The technology could be replicated well enough, was his point; the years of operating it at scale could not be, because the only way to acquire those years is to spend them.

I worked at AWS, and that line gets repeated inside the company the way company lines do, which is often and not always by people thinking hard about it. It is worth thinking hard about, because it is Arrow's argument arriving from the opposite direction — an operator explaining a competitive moat rather than an economist explaining productivity growth, both landing on the same conclusion, which is that experience is an input you cannot buy and cannot rush.

And there is a small irony in where it comes from. Chapter 5 borrowed the one-way and two-way doors from the same company. So the firm that supplied manufacturers with their favourite vocabulary for being careful about irreversible bets also supplied the clearest argument for why the slow accumulating stuff is where the advantage actually sits. The first half has been imported enthusiastically. The second half has not.

Here is what I want you to notice about it.

The learning curve does not pay out in events. It pays out in repetitions. Its returns are a function of how many times you have done the thing, which means value arrives continuously, in quantities too small to see individually, and becomes visible only in aggregate and only in retrospect. It is the erosion from Chapter 5 expressed as an economic law: the grinding is where the shape comes from.

And the machinery we spent the last two chapters describing cannot evaluate anything of that shape. Hurdle rates want a return on a commitment. Business cases want a benefit at a date. A gate is an instrument for judging discrete proposals at discrete moments, and the compounding increment does not present itself as a discrete proposal at a discrete moment. It presents itself as a small thing that will be worth a lot if you run it for two years, which is not an argument the form has a field for.

So the most reliable engine of improvement in manufacturing pays in a currency the approval process does not accept.

. . .

Where the learning actually happens

The learning curve tells you that the improvement pays. It does not tell you where it comes from, and that turns out to matter more.

In 1997 Marcie Tyre and Eric von Hippel published a study in *Organization Science* of what actually happens when new production equipment is introduced into a factory — the messy early period when the machine is installed and nothing quite works. They followed the engineers doing the problem-solving, and what they found was that the process was *situated*. The clues needed to diagnose a problem were physically located in particular settings: at the machine, on the floor, in the hands of the operator running it. Engineers moved back and forth between the lab and the line, not because they were being thorough, but because the information required to solve the problem existed in more than one place and could not be assembled anywhere else.

The finding, stated plainly, is that adaptive learning happens where the work happens, and it cannot be relocated to a more convenient venue.

Now hold that against the geography of a decision. The approval sits in a meeting room, or a portal, or a council that convenes monthly. The learning sits at the line. The person who has the clues is standing in front of the machine and the person with the authority is two floors up looking at a template, and the thing being asked for — a description, in advance, of what the improvement will be worth — is precisely the thing that only becomes knowable in the setting the decision is not taken in.

Maria's rebuild is a small instance of exactly this. Nothing in her fourteen years of experience told her to cluster by failure mode instead of by supplier. No workshop would have produced it. It came out of a mistake made in the sixth week of running the thing, in the presence of the actual data, and it is worth more than the original design.

. . .

The part that cannot be written down

There is a second reason the form defeats her, and it is older and better known than the first.

Ikujiro Nonaka's 1994 account of organisational knowledge creation rests on a distinction between knowledge that has been made explicit – written, codified, transferable – and knowledge that remains tacit, held in a person, expressed in how they work rather than in any artefact. His argument is that organisations create knowledge by converting one into the other, and that the tacit-to-explicit conversion is both the most valuable step and by far the hardest.

Maria's plant runs on tacit knowledge and it does not know it. She knows the seal supplier's quality dips every August. She knows the third shift logs deviations differently to the first. She knows the downtime on line two has been logged against the wrong press since a system migration nobody ever fully unpicked. None of that is in a system. It is in her, and in a handful of colleagues, and it is the reason her tool works at all – she built the clustering rules around distinctions that only exist in her head.

A business case is an explicit artefact. It is, by construction, a request to convert tacit knowledge into codified claims before the conversion has happened. And the increments most worth having are the ones densest in tacit content, because that is what makes them local, and specific, and effective, and hard for anyone else to have thought of.

So the gate applies its heaviest documentation requirement precisely where documentation is least possible.

. . .

The paradox, and its name

Put the three findings together, because separately they are unremarkable and together they close a trap.

The learning curve says the value of an increment accrues through cumulative running. Situated learning says the design only becomes correct at the site of use, through iterations you cannot specify in advance. Tacit knowledge says the substance of the improvement resists codification. Three well-established results, from three literatures that rarely speak to each other, and all three say the same thing about timing.

The value of a good increment does not exist at the moment it must be justified. It is produced by the thing running, in the place it runs, over a period, and the largest part of it is usually something nobody predicted.

This book calls that the Subterranean ROI Paradox: *the increments most worth having are the ones least able to justify themselves in advance*. Not because they are small. Because their value has not happened yet, and the gate wants it now.

That name is a framing device, not a finding. Wright and Arrow did not write about approval processes. Tyre and von Hippel were studying equipment start-up, not governance. What each of them documented is real and I have cited it; the assembly is the book's, and you should hold it as an argument rather than as a result.

But notice that it retroactively explains something we left unexplained in Chapter 1.

The bootlegging research found that hidden projects, on average, come out *better* than sanctioned ones, and the mechanism proposed for that was delayed assessment — that being hidden buys a project time before anyone judges it. That always sounded slightly like an excuse, a nice way of saying that unsupervised people enjoy themselves more. It isn't. Delayed

assessment works because early assessment is assessment of a thing that has not yet become what it is going to be. The submarine is not protecting the idea from criticism. It is protecting the idea from being measured before there is anything to measure, which is a different and much better defence.

. . .

What Toyota knew

I need to close off a reading of this chapter that would be convenient and wrong.

Nothing above is an argument against structure, or process, or formality. The easy version of the argument says bureaucracy kills increments and the answer is less bureaucracy, and there is a famous counterexample sitting in the middle of the industry this book is about.

In 1999 Paul Adler, Barbara Goldoftas and David Levine published a case study of model changeovers at NUMMI, the Toyota-run plant in California, in *Organization Science*. What they describe is a factory that was intensely bureaucratic — standardised work specified to a degree that visitors found startling, formal routines everywhere — and simultaneously one of the most flexible plants anyone had studied. It did both. The mechanisms they identify are unglamorous: metaroutines, routines for changing routines; partitioning, so that the work of improving is separated from the work of producing rather than crushed by it; switching between modes; enrichment of the production role itself. Underneath all of it, sustained investment in training and a level of trust that took years to build.

Toyota's system is a machine for capturing increments. It is more formal than the system Maria works in, not less. So formality is not the variable.

What is different is when and where the judgement happens. In the Toyota case the change is made at the line, by the people at the line, and then evaluated by what it does — with the standardisation arriving afterwards, as the way a proven improvement is locked in and spread. The sequence runs try, observe, standardise. Maria's plant runs justify, approve, try, and it does that because, as Chapter 5 showed, it learned its decision-making on objects where trying first is not an option. You do not pilot a body shop.

Both sequences are coherent. Only one of them can handle a thing whose value appears in month four.

. . .

The arithmetic nobody ran

Let me put a number next to this, with a warning attached: what follows is an illustration, not a study, and I am not going to dress it up as one.

Maria's tool cost six hours and eighteen euros a month. It returns something like a day a week of a quality engineer, plus one early defect catch in eight months of the kind that keeps bad parts off a line. Run that for two years and the arithmetic is not close. It is not close by two orders of magnitude, and you could halve every figure I have given and it would still not be close.

The Quality Systems programme has eighteen months, two consultancies, a budget Maria found hard to repeat out loud, and a roadmap slide that is very good. It has shipped nothing to her plant.

The point of that comparison is not that the small thing beats the big thing. Sometimes it doesn't, and Chapter 12 will make the case for big programmes properly, because they are the right

instrument for some problems and this book is not an argument for a company made entirely of Marias.

The point is that the comparison was never available. To be compared, both sides have to produce a number in the same units at the same time, and only one of them can. The programme can generate a forecast on demand — a large, confident, thoroughly formatted forecast, produced by people whose profession is producing them. Maria's tool can generate a real figure only after two years of running, and it will never run for two years, because to run at all it needed a forecast in month zero.

So a process that requires estimates in advance does not select for the most valuable improvements. It selects for the improvements whose value is *estimable in advance*, which is a different quality altogether, and one that correlates with size, precedent and distance from the actual work. There is no malice anywhere in this. Every individual step is defensible. The system just quietly optimises for forecastability and calls it rigour.

. . .

Where the evidence stops

Four things in this chapter run further than the sources do, and one of them could sink the argument.

The learning curve is common, not universal. It is not a law of physics and it is not automatic. It shows up where an organisation has the willingness and the ability to learn, and where somebody invests in the learning; where those are absent, cumulative output produces cumulative output and nothing else. It also has a boundary that matters here: the classic learning-curve effect describes improvement *within* a production method. It does not describe changing the method. If your improvement needs a different machine, a different line layout

or a different process altogether, then it is a capital decision, the one-way-door caution from Chapter 5 is the correct instrument, and this chapter does not apply to you.

Some increments cannot be separated from the big bet at all. You cannot half-retool a line. There are improvements that are only available as part of a large indivisible commitment, and for those the gate is not an obstacle, it is the only mechanism that could possibly work. The claim I am making is narrower than “the gate is wrong”: it is that reversible, low-cost, locally-learned improvements are being routed through a process built for irreversible, high-cost, centrally-decided ones, and that the routing is the error.

The empirical base for the situated and tacit parts is two case studies. Tyre and von Hippel studied equipment start-up in a small number of plants. Adler and colleagues studied one factory, and it was an exceptional one — NUMMI is famous precisely because it was unusual, which is a limitation as well as a recommendation. These are the best manufacturing-floor anchors I have, and they are two case studies, and I would rather say so than let the word “research” do work it hasn’t earned.

And here is the study that would prove me wrong. Take a large organisation’s completed small process improvements, several hundred of them, and compare what each one’s pre-implementation forecast said it would deliver against what it actually delivered. If the forecasts turn out to be decent — if the benefit an improvement was predicted to produce broadly matches the benefit it produced — then the paradox in this chapter dissolves, the gate is doing exactly what it is supposed to do, and the increments that die at it deserved to. My expectation is that forecasts will systematically miss the largest benefits, because the largest benefits are the ones discovered in use. But I would like to see it done, and as far as I can find, at scale, it hasn’t been.

The trap, complete

Part II is finished, and it is worth saying out loud what it has assembled, because the pieces only become alarming in combination.

Chapter 4 showed that an increment can be killed without ever being judged on merit, because it only has to land on somebody else's territory. Chapter 5 showed that it is rarely even judged, because capital, attention and the calendar all bend toward mass, and a six-hour tool has no mass to bend them with. And this chapter has shown that the thing being buried is not sentimental. Small compounding improvement is the best-documented engine of productivity growth in manufacturing, it has been for ninety years, and the reason it cannot get through the door is not that it is worth too little. It is that it cannot prove what it is worth until it has been allowed through the door.

The trap is now fully visible from both sides. The organisation cannot fund what it cannot forecast. The increment cannot forecast what it has not yet done. And the person holding it — who can see all of this perfectly well without any of the vocabulary in this book — draws the obvious conclusion, which is to stop asking and carry on quietly.

Which leads to an answer that will feel, right now, completely satisfying. If the increment does better underwater, and the gate can only damage it, then leave it there. Protect the submarine. Let the fleet sail.

Hold that conclusion for as long as you can, because Part III is going to take it apart.

Surface Check

So — *should Maria raise her hand on Tuesday?*

The case for her tool is now stronger than it has ever been — ninety years of evidence say the small compounding thing is where the returns actually live — and she still cannot make that case in the room. Everything that makes it worth having appeared after month zero, and the room will ask her what it was worth in month zero. She has already tried to answer that question, alone, on a Tuesday evening, and could not.

So: not yet. But the reason has moved. It is no longer that her idea is weak or that the room is hostile. It is that she is being asked for a proof that cannot exist yet, and the only way to produce it is to keep running the thing where nobody is asking.

INSTRUMENT

The Compounding vs. Big-Bet Calculator

Less a calculator than a comparison you have almost certainly never been shown, in four parts. It takes an hour and you will need someone in finance to be slightly indiscreet.

Part one. Find the forecast.

Take the largest completed programme in your area — completed, not running, so the answer exists. Find its original business case and read the expected benefits. Not the slide. The document, with the numbers that got it funded.

Part two. Find out what happened.

Now establish what it actually delivered, in the same units. This is the hard part and the fact that it is hard is itself the finding. In most organisations the forecast is preserved with great care and the outcome is never assembled at all, because by the time the

outcome is knowable the programme has been closed, the sponsor has moved and the next business case is already in draft.

Write down the gap. Keep it in front of you for part four.

Part three. Now try to do the same for an increment.

Pick one small improvement someone in your area actually made — a changed sort order, a script, a checklist, a tool built in an afternoon. Look for its business case.

There isn't one. Either it never went through a process, or it went through and died. That absence is the second finding: your organisation demands forecasts from the class of improvement that cannot produce them, and does not check the forecasts from the class that can.

For that improvement, do the arithmetic the process never asked for. Not a lump-sum benefit — a rate. How many times per year does it run, what does one run save, and what is that over three years? Then set against it what it cost to try, and what it would cost to stop.

Part four. The question this is all for.

Go back to the increment and ask: *what would this have needed to promise, in writing, in advance, to clear our process?*

Then ask whether anyone could have promised that in good faith before building it.

If the answer to the second question is no — and for the improvements worth having it usually is — then your process is not a filter on quality. It is a filter on predictability, and it has been quietly running for years, and every good increment it turned away is still out there somewhere, running on somebody's personal account, saving them a day a week, and leaving with them when they go.

Next: Part III. Why it's never been fixed — starting not with a mechanism, but with a graveyard.

Maria is still sitting at her desk. The meeting is Tuesday.

PART III

Why It's Never Been Fixed

Part II described a machine, and a machine is a comfortable thing to be handed, because machines can be re-engineered and everyone in the room knows somebody who does that for a living.

Part III is about why nobody has. Sixty years of documented evidence, a problem most executives recognise within a sentence of hearing it described, a remedy that is not technically difficult — and it is still here. That combination ought to be more alarming than it usually is.

Three chapters, and they move steadily away from systems and toward people. The first looks at the programmes built to solve exactly this and asks what killed them. The second asks who gets promoted, and what that selects for. The third asks what all of it does to the person, which turns out not to be a feeling.

This is the least comfortable part of the book.

The Graveyard of Good Programs

In month seven, Maria went looking for a legitimate route.

She had not decided to take one. That distinction keeps mattering. What she had was an arithmetic problem left over from the previous chapter: her tool is worth what it is worth only if it runs for years, because a day a week is a rate and the seal defect was a single event in eight months, and neither of those means anything over a quarter. Two years is the minimum honest window. Three would be better. So the question she was actually asking was not whether the plant would let her do this. It was whether whatever let her would still exist when the arithmetic came in.

She spent about forty minutes on the intranet. This is what she found.

The suggestion scheme is alive. Actually alive, in the sense that a form exists, a committee reviews submissions quarterly, and there is a payout with a cap that was set in a currency and a decade that no longer apply. She knows two people who have used it. One of them got a hundred and fifty euros for a fixture change that saved a shift's worth of rework a month, which he still mentions, and which she cannot decide is a good story or a bad one. The form has a mandatory category field. The

categories are Safety, Quality, Cost, Ergonomics and Environment. Below the category field there is a field labelled *attach drawing or photograph*.

There is no category for having built something.

The idea portal is two and a half years old. It launched with an app, and a video from somebody at head office standing in front of a wall of screens, and a competition. Maria logged in – the login still works, which surprised her – and the news feed on the landing page has one item on it, dated nineteen months ago, announcing the winners of the competition. Nobody has removed the portal. Removing the portal would be somebody's job, and it is not anybody's job.

And then there was Improvement Time, which is the one she actually remembered, and the one worth the paragraphs.

Improvement Time was two protected hours every second Friday, with no reporting requirement attached, which was the part that made it unusual and the part that made it good. The plant manager at the time launched it himself and came to the first session. Maria went to three. The first one was better than she expected – eight people, a flipchart, and an hour in which somebody from logistics said something about pallet labelling that she is still using. The second had four people, because the first shift had run long. The third was cancelled on the morning, because a supplier escalation came in overnight and everyone with a quality badge was needed in a room on the other side of the plant.

Which was correct. That is the thing to hold on to. Cancelling was the right call that morning, and it was the right call the next time, and the time after that, and nobody ever made a wrong decision anywhere in this story.

There was no fourth session. There was no announcement that there would be no fourth session. At some point in the following months Improvement Time stopped being a thing that

was happening and became a thing that used to happen, and Maria could not tell you the month it changed, because nothing changed in any month. The plant manager moved to another site about a year later. She does not think those two facts were connected at the time, and she has since revised that view.

The poster is still up. It is in the corridor by the coffee machine, on the way to the room where Tuesday's meeting will happen, laminated, with one corner lifting away from the wall. It has a QR code on it. The QR code still resolves, to a login page for a system she does not have access to.

Maria was not angry about any of this, and sitting there she noticed that she had never been angry about it, not once, in four years of walking past that poster. It is difficult to be angry at something that stopped. There is no decision to object to, nobody did anything to her, and the only evidence that anything was ever taken away is a piece of laminated card that nobody has got round to removing.

What she registered instead was the shape of the answer. She had gone looking for a route that would last three years. The longest-lived object she found in forty minutes of searching was the poster.

Part II ended by telling you to protect the submarine. Protection is a programme. This chapter is about what happens to programmes.

. . .

The headstones are good

I want to be clear about what is in the graveyard, because there is a lazy reading of this chapter in which the fixes died of being bad.

They were not bad. Several of them were better than anything that has been written about the problem since.

3M has had a version of its fifteen per cent since the late 1940s. Google offered twenty per cent from around 2000. Kaizen and the continuous-improvement family have been deployed in tens of thousands of plants for forty years, with training budgets, certified practitioners and a literature. Suggestion schemes are older than all of it. And the single most interesting intervention in this whole book, which we will come to properly in Chapter 10, is Bernstein's curtain — the one design that produced a measured, sustained improvement in a real factory, and it did that by giving managers less visibility rather than more.

Every one of these was designed by competent people who had correctly identified the problem. That is what makes the graveyard interesting. Chapter 2 registered the instinct — companies noticed that unofficial work produced value and tried to authorise it — and promised you we would watch it fail repeatedly. This is the chapter where I make good on that, and the failure mode is not what most people assume. The programmes did not get rejected. They got launched, and then they thinned out, and then at some point they were being described in the past tense by people who could not name the day.

. . . .

The most famous fix

The Google case is the one everybody knows, and it is worth going through carefully, with a warning attached: what follows is journalism and memoir. It is a case narrative, not a study, and I will flag where it thins.

Twenty per cent time ran from around the turn of the century. Engineers could spend a day a week on work of their own choosing, without seeking approval for the choice.

Marissa Mayer, then a vice-president there, used to tell a story about where Google News came from. After September 11 a researcher named Krishna Bharat was going to a dozen or more news sites a day trying to follow the story, got tired of it, and wrote a crawler to cluster the articles for him. He emailed it around. Somebody who received it looked at it and thought it was not a private tool but a product.

Read that again with the previous six chapters in mind. A person automating their own Friday afternoon, for themselves, with no product intention anywhere in the account, producing something the company had not asked for and could not have specified. That is Maria with a different job title, and the only reason it became a product instead of a private habit is that the environment it happened in did not require it to be justified in advance.

Mayer also supplied the number that made twenty per cent time famous. In a talk at Stanford she said she had mapped the product and feature launches of the second half of 2005 and worked out which came from twenty per cent time and which came through the normal process, and the answer was half.

Take that figure the way it deserves to be taken. It comes from a talk, given by an officer of the company, about her own employer's most celebrated practice, and it is her own assessment rather than an audited one. It has been repeated ever since in books and slide decks by people who did not check. By the standard this book set in Chapter 2, it is a claim, and I am reporting it as a claim.

But now notice what happened next, because the claim's accuracy turns out not to matter much.

By 2013 the programme was being written about in the past tense. The reported causes are unglamorous and they are all recognisable from Part II. Post-IPO quarterly pressure, and the reporting rhythm that comes with it. No formal structure and no

tracking, so nothing to point at. Uneven adoption — the figure that gets repeated is something like one engineer in ten using it consistently, which is also a practitioner number and also unaudited. And a steady re-centralisation of innovation into elite labs, which is Chapter 5's big factory reasserting itself in a company that was supposed to be immune to it. Laszlo Bock, who ran people operations there, described the thing as having waxed and waned with the level of leadership support.

There was no cancellation. That is the part I find hardest to get past. Twenty per cent time was not ended by a decision anybody can date. It became harder to take, then it required a manager's agreement, then engineers started using the phrase "a hundred and twenty per cent time" about it, which is a programme's obituary written by the people inside it.

And the death is contested even now. In 2015 Mayer said the twenty per cent had never really been twenty per cent — that it was a full workload plus whatever else you wanted to do. A Google engineer publicly disagreed, pointing at his own long-running twenty per cent project and saying he was not working a hundred and twenty per cent of anything. Both of them were there. Neither of them is lying.

So here is the summary of the most famous fix in the history of corporate innovation. It ran for over a decade, it was credited by a company officer with half of a year's launches, and it declined without an announcement, in a way that leaves the people who were present disagreeing about whether it existed as described and about when it stopped.

A programme with that record could not defend its own budget line. I think that is the single most important fact in Part III, and the reason for it is sitting in the chapter you just read.

Ask what twenty per cent time delivered last year and you get a story about 2001. Ask what it will deliver next year and nobody can answer, because the whole design principle is that nobody is

specifying the work in advance. The programme's returns arrive late, arrive as something else, and get attributed to whoever shipped the product rather than to the slack that produced it. Which means that when the programme comes up for review – and everything comes up for review – it has to justify itself with exactly the forecast it exists to make unnecessary.

The Subterranean ROI Paradox does not only apply to increments. It applies to the mechanisms built to protect increments. The fix is subject to the gate.

. . . .

The manufacturing version, which is better documented and less famous

The Google case is vivid, and vividness is not evidence. The stronger material is in the continuous-improvement literature, where the same thing has been studied for longer by people with no stake in the outcome.

Nicola Bateman's 2005 paper in the *International Journal of Operations & Production Management* is the one whose title says it. She called sustainability the elusive element of process improvement – not achieving the improvement, sustaining it. Analysing forty improvement activities, she found five conditions associated with activities that actually landed, and five more associated with the smaller number that went on to become continuous. In 2014 Rachel McLean and Jiju Antony published a systematic review in the *International Journal of Productivity and Performance Management* on why continuous-improvement initiatives fail in manufacturing environments, which is a topic with enough failures behind it to warrant a systematic review.

The mechanism these converge on is simple enough to state in a sentence, and I have not found a better sentence for what happens to fixes in general.

The conditions that enable an improvement are temporary. The conditions that generate the problem are permanent.

Everything on the enabling side has an end date. The consultant leaves. The workshop week finishes. The sponsor is promoted or transferred. The training budget covers this year. The attention that made the whole thing feel important lasts about as long as attention lasts, which Chapter 5 already told us is not long and is not yours to allocate.

Nothing on the generating side has an end date. The volume targets are permanent. The escalations are permanent. The headcount is what it is, the schedule is the schedule, and the thing that made the process bad in the first place — whatever combination of history, cost pressure and organisational geometry produced it — is all still there on the Monday after the programme finishes.

So the permanent conditions outlast the improvement rather than defeating it, which is cheaper and requires nobody anywhere to oppose anything.

Session three of Improvement Time was cancelled because of a supplier escalation. That is the whole mechanism in one Friday morning. The escalation was permanent in the sense that matters here — not that one escalation recurs, but that there is always an escalation, and the two hours were always going to be the thing that yielded.

There is a set of figures that circulates about how fast this decay happens. You will have seen them: some large percentage of improvements gone within twelve or eighteen months. I am not going to repeat them, for the same reason Chapter 2 refused the famous innovation-failure-rate numbers. They are practitioner and consultant figures, they get cited to each other

in a loop, and when you follow them back there is nothing at the end. The direction is documented. The magnitude is folklore, and this book does not launder folklore into fact even when the folklore is pointing the way I want to go.

. . .

What a programme actually is, on a cost sheet

Underneath all of these fixes, from 3M's fifteen per cent to two protected hours in a plant nobody has heard of, there is one resource being spent, and it has a name in the literature.

Cyert and March called it organisational slack in 1963: resources beyond what current operations strictly require. Time not committed, attention not booked, budget not allocated to a named deliverable. Slack is what the underground has always run on, and every programme in this chapter is an attempt to create some deliberately and hand it to people.

Nohria and Gulati studied the relationship between slack and innovation and found an inverted U, which is the honest shape. Too little starves experimentation and there is nothing to divert. Too much produces undisciplined projects that should have been stopped. There is a middle, and the middle is where you want to be, and I have never once seen a large manufacturer worry that it might be to the right of it.

Here is the problem with slack, and I think it is the mechanism that explains the graveyard better than anything else I have.

The cost of a protected-time programme is legible. It is hours, it converts to headcount equivalent, it lands under a budget code, and a controller can produce the figure in an afternoon. Two hours a fortnight across four hundred people is a number, and it is not a small number, and it will be on a slide at some point.

The benefit is not legible. It arrives late, it arrives as something else, it gets credited to the person who shipped the thing rather than to the arrangement that let them build it, and a great deal of it never surfaces at all because it stays as somebody's private improvement to their own Friday. There is no line for it anywhere.

So the fix appears on the cost side of the ledger with complete precision and on the benefit side not at all. In a good year nobody looks. In a bad quarter somebody looks at everything, and this is the cheapest thing in the building to stop, because stopping it requires no negotiation with a customer, no contract to unwind, no capital to write off, and no one person whose job disappears. You stop it by not scheduling it.

And there is a real conflict here rather than a stupidity, which is worth saying because this book keeps insisting that every step is defensible. Forty years of improvement doctrine in manufacturing has been about the removal of slack. That doctrine was right about a great many things and it built the modern plant. But it cannot distinguish between slack that is waste and slack that is the input to the learning curve, because the two look identical at the moment you measure them. They are only distinguishable afterwards, and Chapter 6 has already explained why afterwards is too late to be useful.

Which gives the programme its half-life. A fix of this kind survives as long as somebody with standing is prepared to defend a cost line with no matching benefit line, on the basis of a belief they cannot evidence. That is a personal act, not an institutional one. So the programme's lifespan is roughly the tenure of the person who believes in it, and Maria's plant manager moved sites about a year after Improvement Time stopped happening, and I do not think she is wrong to have revised her view about the connection.

The fix that eats itself

There is a second way these programmes die, and it is the one that catches the organisations that take the problem seriously.

Recall the finding from Chapter 1 that Chapter 6 finally explained. Criscuolo, Salter and Ter Wal found that bootleg efforts were associated with higher individual innovative performance, and the mechanism was delayed assessment — the hidden project gets to become what it is going to be before anybody judges it.

That is not the only thing they found. In the same study, during periods of R&D process formalisation, more bootlegging was associated with *less* innovation. The behaviour did not change. What changed was the surroundings, and the same behaviour in formalised surroundings stopped producing the same result.

Now think about what happens to a well-run protected-time programme over its first two years, in an organisation with competent management, which is to say an organisation where things get managed.

Somebody sensibly asks how many people are using it, so there is a sign-up. Somebody sensibly asks what has come out of it, so there is a showcase each quarter. Somebody sensibly asks how we know the time is being used well, so there is a field on a timesheet. Somebody sensibly asks whether we should be steering it toward the strategic priorities, and now there is a list of suggested themes, and the list is not compulsory, and everybody understands the list.

Each of those is a reasonable request from a reasonable person. Together they have rebuilt the gate inside the sanctuary. A protected day with a sign-up, a tracking field, a quarterly

showcase and a themes list is an approval process with better manners, and the thing that made the underground work — that nobody is assessing the embryo — has been removed by people trying to look after it.

So the programme has two available deaths and it must choose one. Audit it, and it stops being a nursery, because assessment is the thing it existed to defer. Leave it unaudited, and it has no evidence of its own value when the bad quarter arrives, which is the death described in the previous section.

And notice what actually did the killing, because nobody judged the programme and found it wanting. The looking was the whole event, and what the programme became afterwards was set by the looking rather than revealed by it.

I have not found a version of this trap that anybody has solved in the general case. Bernstein's curtain is the closest thing to an escape I know of, and it works precisely because it does the opposite of what a manager's instinct demands: it reduces what can be seen, on purpose, and then measures the line rather than the people. That is Chapter 10, and it is the reason that study matters more than anything else in this book.

. . .

Where the evidence stops

Four limits, and the first one is about the chapter's most quotable material.

The Google case is journalism, and its death is disputed.

Everything in that section comes from press coverage, from talks given by executives about their own company, and from memoir. There is no study. The adoption figure and the launch-share figure are both unaudited claims that have been repeated into furniture. And the participants do not agree on the basic facts: one senior figure has said the twenty per cent was always

effectively a hundred and twenty, and engineers who were running twenty per cent projects at the time have said publicly that this is not what it was like for them. I am using it anyway, because it is the most legible instance of a pattern the manufacturing literature documents far more carefully and far less vividly. But the careful documentation is Bateman and McLean, not Google, and if you take one citation out of this chapter take one of those.

Programmes do survive, and I cannot fully explain the ones that do. 3M's fifteen per cent has lasted more than seventy years, through every pressure described above and several recessions. So the claim here is a tendency, not a law, and any honest version has to account for the exception. My guess — and it is a guess — is that it survived because it predates the reporting culture that would have made it justify itself, and because it long ago stopped being a programme and became something the company understands itself to be, which is a much harder thing to remove in a bad quarter than a line item. I have no test for that and I am not going to pretend otherwise.

The decay figures are folklore. Stated above; restated here because it is the kind of number that gets reinserted by a helpful reader. The direction of travel is documented. The percentages are not.

And here is the study that would settle it. Take protected-time and continuous-improvement programmes across a few hundred firms, with dated launches, and measure how long each survived — against two events specifically: a change of sponsor, and two consecutive bad quarters. Then compare that survival against matched organisational initiatives of similar annual cost whose benefits *are* attributable to them. If the two groups die at the same rate, then this chapter has discovered only that initiatives end, which is true, well known and not worth a chapter. My expectation is that the unattributable-benefit

programmes die substantially faster, and that sponsor tenure predicts a programme's lifespan better than any measure of what the programme actually produced. As far as I can find, nobody has assembled that dataset, and it would not be hard to assemble.

. . .

What this does to the answer you were just given

Part II ended on a conclusion that felt complete. If the increment does better underwater and the gate can only damage it, then leave it there, protect it, let it run. I asked you to hold that as long as you could.

Here is the first thing wrong with it.

Protection is a programme. Whatever form it takes — protected hours, a sanctioned sandbox, an exemption, a tolerant manager who has decided not to ask questions — it is an arrangement that costs something legible and returns something illegible, and it therefore has a half-life. The graveyard is full of the good ones. And the payback period of the thing being protected is longer than the half-life of the protection, which is the whole problem in one sentence: Maria's tool needs three years to prove what it is worth, and nothing in that building has lasted three years except a laminated poster.

There is a second thing wrong with it, and it is worse, and this chapter cannot pay it off.

Look again at how these programmes ended. Nobody cancelled twenty per cent time. Nobody cancelled Improvement Time. There was no meeting, no decision, no record, and no one to be angry at — the thing simply thinned out until it was in the past tense. Hold that shape in your mind, because it is going to come back in Chapter 9 wearing different clothes, and when it does you will find that a protection which is quietly withdrawn

does more damage than a protection that was never offered. Everyone who used it learns something from the withdrawal. The learning is the expensive part, and it is not on any ledger.

So Part III opens with the fixes already existing and already failing, and the obvious next question is why. The proximate mechanisms are above and I think they are real. But there is a harder version of the question underneath them, which is this: these programmes get launched by senior people who understood the problem well enough to launch them, and they get allowed to lapse by senior people too, often the same ones, often without anybody noticing they have changed their mind.

The answer is not in the design of the programmes. It is in who is doing the deciding, and how those people got to be the ones deciding.

. . .

Surface Check

So – should Maria raise her hand on Tuesday?

The question has changed shape. She is no longer only asking whether the room will say yes. She is asking whether the yes would still be there in two years, because two years is when her tool's arithmetic finally comes in – and she has just spent forty minutes establishing that nothing in that plant, including the mechanisms built specifically for this, has ever lasted that long.

So: not yet, and for a new reason. What she distrusts now is the durability of whatever sits on the other side of the gate.

INSTRUMENT

The Program Half-Life Check

An hour, and unlike the earlier instruments this one is mostly archaeology. You are going to find out what your organisation has already tried, which in most places nobody knows, because the people who remember have moved and the record was never kept.

Part one. Count the graves.

List every programme launched in your area in the last ten years whose purpose was to get more small improvements out of the people doing the work. Suggestion schemes, idea portals, protected time, innovation days, hackathons, improvement circles, kaizen waves, digitalisation initiatives, whatever the local name was. Include the ones you inherited and never used.

For each, write down four things: the year it launched, who sponsored it, the date of the last real activity you can evidence, and — this is the one — **who decided to end it**.

Expect that last column to be almost entirely empty. That absence is the first finding, and it is not an administrative oversight. A programme that nobody decided to end is a programme nobody had to defend ending, which means nobody had to say out loud what was being given up.

Part two. Work out the half-life.

For each dead programme, take the months from launch to last real activity. Now put next to it the tenure of its sponsor in the role.

If those two columns look alike, you have the mechanism, and you also have a forecast. Whatever you launch next will last about as long as you stay, and no longer, unless something about it is different from everything in the column above it.

Part three. Find it on the cost sheet.

Take the programme running right now — there is usually one — and answer two questions.

Where does its cost appear? Budget code, headcount equivalent, hours per person per period. Somebody can give you this figure today.

Where does its benefit appear? Not the anecdote. The line. Which reported number would move, by how much, and who would attribute the movement to this programme rather than to the person or team who shipped the thing?

If the cost has a line and the benefit does not, you already know what happens in the first bad quarter, and so does everyone who has been here longer than you.

Part four. The three questions, asked before you need them.

Ask these of the programme you are running, or the one you are about to launch:

1. Does it survive you leaving? Name the person who would defend it if you were not in the building. If you cannot name them, its half-life is your notice period.
2. Does it survive two consecutive bad quarters? Not whether it *should*. Whether it *would*, given who will be in the room and what else is on the list.
3. Does it survive being measured? Work out what tracking you would add under pressure to justify it — the sign-up, the timesheet field, the quarterly showcase — and ask whether the thing still does its job once that tracking exists. If it does not, you are choosing between two deaths and you should choose deliberately rather than by drift.

If the answer to any of these is no, you have a choice, and only one of the options is honest. You can launch it anyway and say nothing, which is the standard practice. Or you can tell people at the start what its expected lifespan is, and what you will do for them if it ends before their work is finished.

That second option sounds like a weak way to launch something. I think it is the only version that does not cost you more than it delivers, and Chapter 9 explains why.

The poster is still on the wall by the coffee machine.

Next: Chapter 8 — why the fixes die, which turns out to be a question about who gets promoted.

Maria is still sitting at her desk. The meeting is Tuesday.

CHAPTER EIGHT

Promoted for the Big Bet

In month eight, Maria read the attendee list. She had not looked at it when the invitation arrived, because nobody reads attendee lists, and she only opened it on the Thursday before for no better reason than that she was avoiding something else. There were fourteen names. Most of them she could place — her manager, two other quality leads, somebody from IT whose name she had seen on ticket responses and never in a room, three people from the Quality Systems programme, which she noted without feeling much, having sat through three of their update presentations and knowing exactly what the roadmap slide looks like. And two-thirds of the way down, a name she has known for a long time.

He joined the year before she did, same function, same grade, same building, and for the first few years they were on more or less the same list of people. He is now two levels above her, in a role with the word *excellence* in the title, and he will be in the room on Tuesday. Maria can tell you the exact thing that put him there, because she was standing next to it when it happened.

Seven years ago the scrap figure on line four was wrong, and it had been wrong since a system migration, in a boring and specific way: rework was counted twice, once against the station

that found it and once against the station that caused it, so line four had been carrying a scrap number about a third higher than reality for two years. Nobody had caught it, because nobody questions a number that has always looked like that. It looks like a property of line four. Maria caught it the way she catches most things, by noticing that a figure had not moved when it should have, and she took it to the man who owned the mapping, and he fixed it in an afternoon. That was a Thursday, and she remembers because she had been planning to leave early.

He caught it about three weeks later, on his own, coming at it from a different direction, and this is where the two of them stop being the same person. He did not take it to the mapping owner. He wrote it up, and he wrote it up as what it plainly also was — evidence that the migration had left errors of this kind all over the plant, that nobody had audited them, and that the plant had been making decisions for two years on numbers it had never checked. He proposed a data-integrity workstream and he got one. It ran for eighteen months with a steering group and a name and a colour on a slide, it found eleven more errors, two of which mattered a great deal, and when he presented it at a regional quality meeting he was promoted about four months later.

Maria has thought about this on and off for seven years and has never once concluded that he was wrong, which is the part worth being careful about. He was not exaggerating, he was not empire-building, and he was not taking her fix, because he did not know about her fix — by the time the workstream started, line four had been correct for a month and there was nothing anywhere recording who had made it correct or when. He had identified a real problem across the whole plant, one that Maria had walked straight past. She had fixed one number; he had noticed a class of failure. He was better than her at something, and the something was not engineering.

What he had was a version of the problem with a shape to it, the kind of shape you can hang a sponsor on, and a budget line, and a start date, and a slide with eleven errors on it. Maria's version was fixed on a Thursday afternoon by a man who was going to be at his desk anyway, it produced nothing that anyone could point at afterwards, and there is nowhere in that building where it is written down.

So on Tuesday she will sit in a room with the person that story is about, and she has spent most of this week working out what she actually expects from him. He will not be hostile, because he has never been hostile to her about anything. He is good at his job, he is not stupid about tools, and he has spent two years asking in meetings why nothing moves faster. Which is the problem, laid out with unusual clarity in a corridor on a Thursday evening: if she brings the tool into that room, the outcome she should be afraid of is not that he dislikes it, but that he likes it.

She knows what he does with things he likes.

. . .

The lazy version of this chapter

There is a way to write this chapter that would be popular and that I am not going to write. In that version the reason nothing gets fixed is that the wrong people are in charge — they are vain, they want monuments, they care about their next job more than about the plant, and the small useful thing dies because it flatters nobody. It is a satisfying story, and most people who have worked in a large company have at least one senior figure in mind who fits it exactly.

I think it is mostly wrong, and worse than wrong, it is useless, because it aims the remedy at the wrong thing. If the problem is bad individuals then the answer is better individuals, and every

organisation on earth has already tried that at enormous expense, with leadership development budgets that would fund a small research council.

The claim in this chapter is different and it is harder to live with. You can staff the whole chain of command with capable, decent, well-meaning people — people who would pass any character test you cared to set, who are not vain, who want the plant to do well — and the machinery from Part II still produces the same result, because what gets a person chosen and what makes a person right are two different things, and nothing in the choosing is able to tell them apart. None of that is a claim about anybody's motives. It is a claim about what a promotion decision can see in the hour it has.

. . .

Two questions that look like one question

The finding this chapter rests on is a review of the research rather than a single study, and its power sits entirely in a distinction that sounds like hair-splitting until you hold it for a moment. In 2015 Emily Grijalva and colleagues published a piece in *Personnel Psychology* pulling together decades of work on narcissism and leadership, and they kept two outcomes separate. The first is *emergence*: does this person become a leader, get picked, rise, end up in charge. The second is *effectiveness*: once they are in charge, does it go well.

Narcissism predicts the first and does not predict the second.

The relationship with emergence is positive and it is not small. The relationship with effectiveness is near enough nothing, and once you allow for extraversion, which travels alongside narcissism and does much of the visible work, what is left looks like a curve rather than a line — some is useful, more is not more.

A word about that term, because it does damage every time it appears. What these studies measure is a normal-range personality trait rather than the clinical disorder: confidence in one's own judgement, comfort with authority, a willingness to be the centre of a room. Plenty of people who score high on it are people you would call impressive. Very few are people you would call narcissists in conversation, and reading this literature as a list of bad people is the fastest way to learn nothing from it.

The result that should worry a selection panel came a few years earlier. Brunell and colleagues in 2008 ran the emergence question in group settings and found the high scorers rising to the top even where the people watching had been trained to assess leadership.

Training the observer did not remove the effect, which tells you the signal is not being missed through carelessness.

Cameron Anderson and colleagues supplied the reason in 2012, in work on what they called the status-enhancement account of overconfidence: overconfident people gained standing in their groups, and they gained it because the overconfidence was read as ability, and it was read as ability because nobody spotted it as overconfidence. The group was not tolerating a flaw it could see. It could not see it.

Tomas Chamorro-Premuzic assembled all of this for a general audience in 2019, and I will flag that book as what it is, a trade book drawing on other people's data rather than a study, because his central line is still the cleanest statement of the mechanism I know: confidence gets used as a stand-in for competence, and the two have very little to do with each other.

Now set that against the room a promotion actually happens in – an hour of conversation, three or four busy people, a written record of achievements the candidate wrote themselves, and perhaps a reputation arriving second or third hand. Confidence is fully visible in an hour. Competence is not visible

in an hour, and in a great many jobs it is not cleanly visible in three years, because the comparison does not exist and nobody ever gets to watch what the plant would have done under the other candidate.

So the panel does what every other mechanism in this book does, and takes the signal it can read in place of the one it cannot, because the one it can read is the only thing in the room.

That should be familiar by now. In Chapter 6 the benefit that could be forecast pushed out the benefit that could only be discovered. In Chapter 7 the cost of protected time was a budget line and the benefit was not, so the programme died in the first bad quarter. This chapter is the same trade in a third currency. It is one mistake showing up in three places, and the currency is never really the point — whatever can be measured easily will stand in for whatever cannot, and there is always something that can be measured easily.

. . .

What the selected people then do

Half of this chapter's job is working out who gets picked. The other half is showing that the picking has consequences for the particular thing this book cares about, and here the evidence is older and much better funded, because it is finance research and finance research follows the money.

Michael Jensen's 1986 paper on free cash flow is where the idea starts that managers over-invest and over-expand for reasons unconnected to the return, and it names the reason plainly: status, power, freedom of action and pay all rise with the size of the thing you run, so growth pays the person deciding whether or not it pays anybody else.

Ulrike Malmendier and Geoffrey Tate then measured it. Their work in 2005 and 2008 picked out overconfident chief executives by their habit of holding company stock options long past the point where a sensible person would have cashed them in, which is an elegant way to catch someone who believes their own forecasts, and found that those executives buy other companies far more often — roughly sixty-five per cent more likely to make an acquisition, more likely to destroy value when they did, and most of all where the firm was sitting on plenty of its own cash and therefore never had to convince an outside lender of anything.

Richard Roll had proposed the shape of this in 1986 as the hubris hypothesis, and Mathew Hayward and Donald Hambrick tested part of it in 1997 across a hundred and six large acquisitions, finding that signs of hubris in the chief executive predicted how much of a premium the company paid. The more certain the person, the more they paid.

Arijit Chatterjee and Hambrick ran the personality version directly in 2007, in a study that gets quoted more for its cleverness than for its result. They could not sit chief executives down with a questionnaire, so they built their measure out of the clues those executives leave behind: how prominent the chief executive's photograph is in the annual report, how often they are named in press releases, how much they say "I" in interviews, and what they are paid next to the second-highest-paid person in the firm. Higher scores went with bolder strategy, grander gestures, and more and larger acquisitions.

The result I want is the last one, and it is the one that keeps this chapter honest. Performance under those executives swung much harder in both directions, bigger wins and bigger losses, and on average came out no better and no worse than under everybody else.

That removes the villain from the story completely. These are not people making bad decisions; they are people making risky ones, and risk is not a moral category. A firm that never takes a big swing is not obviously in better health than one that does.

But risk is not neutral for the person taking it, and this is where the individual and the company come apart. The good outcome of a big visible bet has a name on it, and the name is theirs. The bad outcome spreads out and lands nowhere in particular, because by the time it shows up the market had moved, the integration was harder than anyone expected, the supplier failed, and in any case they are running something else now. Same odds, very different personal payoff. Nobody has to be cynical for that to work, and nobody has to notice it either.

There is one more piece and it arrives with a warning attached, as promised in Chapter 2. Barry Staw's 1976 study, the one everybody knows by its title about being knee-deep in the Big Muddy, found that people who were personally responsible for a failing course of action put more resources into it than people who had inherited the same failing course from somebody else, with self-justification as the driver.

It is among the most cited results in the field and it fits this chapter almost too neatly, which is exactly when a book should slow down, because it has a replication problem. Martens and Orzen went back to it experimentally in 2021 and did not reproduce it — no increase in spending after bad news — and related work has reported the opposite effect, where a larger initial commitment made carrying on less likely. The field evidence for escalation is broad, and the 2012 review by Sleesman and colleagues supports the phenomenon in general, but the famous experiment is disputed.

So it does light work here, which is all I need it for: it is one possible explanation for something Chapter 5 already established on other grounds, that a large programme does not

give back its share of attention when it starts going badly. I am not going to build anything on it. A famous but shaky finding does not get to carry weight in this book just because it points where I was already heading.

. . .

Nobody is promoted for a programme that is still running

Here is the connection the last two chapters have been walking toward, and it is mine rather than anybody's finding, so take it as an argument. A career is built out of events, and that word is carrying more weight than it looks. A launch is an event. A turnaround is an event. A restructuring, an acquisition, a new function stood up, a programme announced with a video from head office — all events, all with a date on them, all attached to a person, and all of them things you can write in the box on the form that asks what you have delivered.

Keeping something alive is a state rather than an event. It has no date and it leaves nothing behind, and the nearest thing to evidence it produces is an absence, which is that the thing did not stop, and an absence cannot go on a slide because the slide would have to say that nothing happened.

So run the sum that a competent, decent, entirely unvain senior person runs without ever putting it into words. There are two options. Launch a new improvement programme, which is visible straight away, belongs to them personally, and gets judged on the launching rather than on the outcome, since the outcome is four years out and they will be somewhere else by then.

Or defend the budget line of the programme their predecessor launched, which is a cost with no matching benefit — Chapter 7 established that much — in a meeting where they would have to argue for something they cannot evidence, on

behalf of somebody who has left, and which if it works will produce results credited to whoever ships the thing.

Both of those are open to a person of perfectly good character. One of them is a career and the other one is a favour.

And this, I think, is the answer Chapter 7 could not give you. Improvement Time did not die because the plant manager stopped believing in it. It died because launching it was the event, the event had already happened, and there was nothing left in it for anybody. He was promoted for the launch, moved sites about a year later, and by then it had already become a thing that used to happen. Nobody in the history of that company has been promoted for a programme that was still running, and nobody has ever been promoted for the absence of a decline.

Which gives the graveyard its shape: those programmes were not killed by the people who launched them, but launched by the same mechanism that then gave nobody any reason to keep them.

The three currencies close here, and they close in one sentence. Money bends toward the big because big decisions cannot be undone, attention bends toward the big because size pulls the room, and careers bend toward the big because a career is a list of things you started and only the big ones count as things you started. Three separate mechanisms with nothing in common, drawn from three bodies of research that have never been introduced to one another, all leaning the same way. That convergence is the argument of this book, and all three legs of it are now on the floor.

. . .

Why the fix has a minimum size

There is a consequence of all this that turns up in every large organisation trying to take innovation seriously, and Chapter 7 walked past it in the Google material without stopping. Twenty per cent time decayed partly through a steady pulling of innovation back into elite units — Chapter 5's big factory reasserting itself inside the one company supposed to be immune to it. With this chapter in hand you can see why it was always going to happen, and the reason is not the one I expected when I started looking.

Start with the leader, sincerely wanting more innovation, whose own contribution will be judged on events. Distributed increments produce no event at all. Four hundred people each improving their own work by some amount nobody measured has no launch date, cannot be photographed, and would be credited to four hundred people, none of them the leader. There is no version of it that reaches a slide except as a claim nobody can check.

A centre is an event. It has an address, a head count, a budget, a name, a director and a ribbon, and it can be shown to the board, to journalists, and to a candidate the company is trying to hire. So the honest version is that nobody chooses the centre over the distributed thing, because only one of the two is on the menu. A leader can pick the centre. There is no procedure by which a leader picks the distributed increment, because there is nothing there to pick — no proposal arrives, no budget is requested, and the thing it would produce has no name.

That much I had before I started writing this section. What I had wrong was the assumption that the centre is simply the wrong instrument, a clumsy tool aimed at delicate work. It is worse than that.

A centre has a cost line. Head count, floor space, tooling, a director's salary, and all of it visible in the plan every year. Chapter 7 established what happens to a cost line with no matching benefit line, so the centre knows from its first week that it will have to show what it is for. Which means it needs work large enough to be worth forty people, because work that is not large enough to be worth forty people is an argument for having fewer of them.

So it acquires a floor. Not a written one — nobody circulates a memo setting the minimum size of an acceptable topic — but a floor all the same, visible in what gets taken on and what gets politely redirected. Below the floor sits every improvement of the kind this book is about: worth a day a week to one engineer, invisible on any reported number, and impossible to describe as a reason for the centre's existence.

Which is the trap from Chapter 7 wearing different clothes. The fix is subject to the gate, and here it is subject to it twice over, because an increment brought to the centre now has to justify itself *and* justify the centre.

Maria's tool cost eighteen euros a month and six hours of a Sunday. It has no threshold to clear, because there is nothing underneath it demanding to be paid for. Move it into a unit with forty people and a floor, and it becomes too small to be interesting to the only part of the organisation officially responsible for being interested in it.

And the distance makes it worse rather than better. Tyre and von Hippel showed in Chapter 6 that this kind of improvement is situated — it emerges where the problem is, from people who have the clues in front of them.

A centre is staffed by people taken off the line, which is what lets it be a centre at all. The same property that makes it visible enough to survive is the property that puts it out of reach of the work it was built to capture.

The same logic explains how a protection that already exists gets quietly starved without anybody starving it. Chapter 7 established that nobody decided to end Improvement Time, and this chapter explains why nobody had to. Ending it took no act at all. Keeping it would have taken a repeated act, every planning round, with nothing coming back the other way.

. . . .

The case for the bold ones

I want to put the other side at its strongest, because it is stronger than this chapter has so far allowed and leaving it out would make the argument cheap. Confidence is not a defect. An organisation facing a real threat needs somebody willing to commit serious resources before the evidence is in, and the person able to do that is precisely the person who scores high on the traits described above. A firm whose leaders only ever acted on complete information would not be careful, it would be dead, and it would die slowly enough that everyone involved could explain why each separate delay had been prudent – which is more or less the thesis of the companion volume to this book.

Chatterjee and Hambrick's average result points the same way. No worse on average, whatever else the big-bet temperament is, means it is not reliably destructive, and a book that read their finding as an indictment would be misreading it.

Adam Grant adds a wrinkle that matters here, with the same trade-book caveat attached. In *Give and Take* he separates people who claim more than they contribute from people who do the reverse, and finds the self-promoters doing well at getting picked and doing well in the short run, while over longer stretches the advantage turns around.

His most useful category is the person who is prickly and difficult and nonetheless acting in other people's interests, who he argues is the most systematically undervalued profile in working life, because the visible signal, which is the manner, points the opposite way to the thing that matters.

Which is this whole chapter again, arriving from the other side. None of it makes the trait the problem. The problem is that the organisation has one way of choosing people, it runs on a signal available in an hour, and it therefore cannot separate four different people who look identical across a table. Which is why the remedy cannot be better people. It has to be a second channel, and Chapter 12 is where that gets built.

. . .

Where the evidence stops

Four limits, and the first is the biggest leap in the book.

Almost all of this evidence is about chief executives, and this chapter is about a plant. Jensen, Malmendier and Tate, Roll, Hayward and Hambrick, Chatterjee and Hambrick — every one of them studies the very top of very large listed companies, mostly through acquisitions, in samples weighted toward finance and technology. Nobody in that literature has looked at a promotion round in a manufacturing company, or a plant manager's spending choice, or the person chairing a use-case review in a factory. The step from *overconfident chief executives buy more companies* to *the man across the table on Tuesday will turn her tool into a programme* is a leap. I think it is a fair one, because the mechanism underneath it — that the reward attaches to visible events with a name on them — is if anything stronger in the middle of a hierarchy than at the top, where there is at least a board and a share price pulling in the other

direction. But it is a leap, it is doing real work here, and if you want to attack this chapter, that is where to attack it.

The evidence about who gets picked and the evidence about what they then do come from two different literatures. Grijalva and Anderson are psychologists measuring who rises, in samples that are largely not senior executives. Malmendier and Chatterjee are finance and strategy researchers measuring what senior executives do. This chapter joins the two: selection favours the trait, the trait produces big visible bets, so the deciders end up tilted. No study makes that join. It is the same kind of assembly as the three-currency argument, and it deserves the same standing — a strong case rather than a proven chain.

Staw is disputed, and the personality measures are stand-ins. The escalation caveat is above. Separately, the clues Chatterjee and Hambrick used — photograph size, counting the word “I”, pay ratios — are ingenious and they are still stand-ins, and it is fair to ask whether a large photograph in an annual report measures a person or measures the house style of a communications department in a particular decade.

And here is the study that would settle it. Take a large organisation with ten years of promotion records. For everyone promoted into a decision-making role in that time, score two things from the record: how many things they launched, and how many things they inherited, kept running and handed on working. Then test which of the two predicts promotion. If keeping things alive predicts advancement at anything like the rate launching does, this chapter is wrong, Chapter 7's graveyard needs another explanation, and Part III loses a leg. My expectation is that launching predicts promotion strongly, that stewardship predicts nothing whatever, and that most organisations will find they cannot even fill in the second column because they never recorded it — and that last part is

not an obstacle to the study but the first result of it. Every large company already holds the data for this. As far as I can find nobody has run it, and it would not require anybody to be asked anything.

. . .

Who she would be surfacing to

Eight chapters in, the room is finally fully described, and the thing to notice about it is that there is nobody in it to be angry with. The gate was reasonable, in Chapter 4. The caution about money was rational, in Chapter 5. The programme's death needed no decision at all, in Chapter 7. And the person at the head of the table was chosen by a process doing its honest best with the only signal it could see in the time it had, and is now behaving exactly as chosen, which is what choosing means. Every step defensible, every step correct from where the person taking it was standing, and the whole thing adds up to a machine that reliably turns a working tool into an eighteen-month programme and reliably fails to notice that it has.

This is why the problem survives being well understood, which is the question Part III exists to answer, because a problem made of bad actors gets solved by removing them and a problem made of reasonable decisions has nothing to remove.

It also puts Maria's Tuesday into its final shape. Her risk is not rejection — rejection would leave her exactly where she is, with a tool that works and eighteen euros a month. Her risk is adoption, that the tool gets taken up by the person best placed to take it up, who will do the only thing his position allows him to do with something he likes, which is to give it a sponsor and a scope and a name, and make it visible enough to be worth his time, and in doing so turn it from a thing that runs on a Sunday into a thing that has to justify a forecast it cannot produce.

She has watched him do exactly this once before, correctly, with a real problem, seven years ago. He was right that time, and that is what makes it so hard to argue with.

There is one thing left, and it is the piece this book has been walking around since Chapter 1. Three chapters have gone on what the organisation does to the idea, and none of them has asked what it does to the person — not on the Tuesday, but over eleven years, and not as a mood but as something she learned and has had no reason since to unlearn. That is Chapter 9, and it is the hardest one.

Surface Check

So — should Maria raise her hand on Tuesday?

The question has stopped being about the room and become a question about one person in it. She has spent seven chapters worrying about a no, and what she has finally worked out is that the no is survivable and the yes is not, because the yes comes from somebody who is good at his job, wants the plant to succeed, and can only take in a small working thing by turning it into a large unproven one.

So: not yet, and now for the reason that is hardest to fix. There is nobody here she needs protecting from. That is exactly the problem.

INSTRUMENT

The Emergence-vs-Effectiveness Check

Forty minutes, a promotion record you can already get at, and one honest column. This one is aimed squarely at anybody who

sits on a panel, and unlike the earlier instruments it asks you nothing about culture and nothing about how anyone feels. Every question below has a factual answer.

Part one. The last five promotions.

List the last five people promoted into decision-making roles in your area, and for each of them write down the single thing they are most associated with – the thing that comes up when their name comes up. Now sort those five things into three piles. A launch: they started something. A turnaround: they took something failing and fixed it. A continuation: they kept something good running that they did not start.

Count the third pile. In most places it is empty, and the empty pile is the finding. That does not mean nobody kept anything alive. It means keeping things alive is not the kind of thing the record holds.

Part two. The missing column.

For each of those same five people, name one thing they inherited, kept alive, and handed on working. You will find this hard, and the difficulty is not about them, so try to work out where you would have gone to look it up. There is no field for it on any form in your building, no section for it in any review template, and no question about it in any interview any of them sat. A thing nobody records is a thing nobody can be promoted for, and a thing nobody can be promoted for is a thing that quietly stops being done.

Part three. Your own two numbers.

Count the things you have launched. Then count the things you inherited that are still running, that carry somebody else's name, and that will outlast you without ever being credited to you. Now look at your last promotion and work out which of the two numbers paid for it.

This is uncomfortable and it is meant to be, though not in the direction you are probably taking it. Nothing here says you were

wrongly chosen. The finding is that neither you nor anyone on that panel could have told the difference from the inside, because the signal that separates the two candidates is not available in the room where the choice gets made. That is a property of the measurement, and you inherited it.

Part four. The panel question.

Before the next selection you sit on, write down in advance, while you are still calm, what evidence available to you in that hour would separate a confident candidate from a competent one. If you cannot name any, then you are about to choose on confidence, and knowing that is the only defence there is.

Then add one question and ask it of everybody: *tell me about something you inherited and made better*. It is a useful question for reasons that have nothing to do with values. A launch story can be told convincingly by the person who launched it. A story about keeping something alive usually cannot be told at all without naming other people, and dates you can check, and a state of affairs that either does or does not still exist, which makes it the one part of the hour that is difficult to perform. You will not fix the selection system with a single question. You will find out how many of your candidates have an answer.

The scrap figure on line four has been correct for seven years. It took an afternoon. There is no record anywhere of who fixed it.

Next: Chapter 9 — what all of this does to the person, which turns out not to be a feeling.

Maria is still sitting at her desk. The meeting is Tuesday.

CHAPTER NINE

The Silence That Follows

Maria was twenty-six the first time, and she has not told this story at work in eleven years.

She had been at the plant three years and she had noticed something about the incoming inspection queue. It ran alphabetically by supplier name, which is how the system had always sorted it, which meant a supplier whose parts stopped the line on Monday got looked at after a supplier whose parts sat in a rack for three weeks. She wrote it up. Two pages, a quarter of queue data, one proposal: sort by line-stop risk instead of by name. No new system. A change to a sort order.

Her team leader thought it was good and sent it up. Somebody above him thought so too, because it went further, and for two weeks Maria heard things — it had been mentioned, someone had asked a question, it might go to the process council.

Then nothing.

Not a rejection. The precision matters here. Nobody said no, nobody said the analysis was wrong, nobody said it was somebody else's decision to make. The thing stopped moving the way a stone stops moving, without any event you could point to.

She asked her team leader after a month. He said he would check. He may have checked.

At three months she asked again and he said something she still remembers: *these things take a while*. He was not being unkind. He was being accurate.

At six months she stopped asking. Nothing happened at six months either. There was no moment. She noticed one afternoon that she had stopped thinking about it, and that the noticing did not upset her much.

The queue is still alphabetical. Maria checked once, a few years ago, and has not checked since.

. . .

What she learned

The comfortable reading is that Maria was discouraged, that discouragement is a feeling, and that feelings pass. On that reading the fix is encouragement and the problem is morale.

That reading is wrong. Maria did not feel something. She learned something, and learning is not a mood. It is an update, and it outlasts the mood that accompanied it.

What she learned is that the relationship between offering an idea and anything happening was too weak to be worth the offering. Not that ideas get punished — hers wasn't. Not that her manager was hostile — he wasn't. Something more corrosive and much harder to argue with: that it does not matter.

And she was right. This is the part to sit with rather than skate past, because the temptation is to treat her conclusion as a distortion, a bit of learned pessimism to be corrected with better communication. It is not a distortion. She observed a system, formed a hypothesis about how it responds to input, ran the test once, and got a clean result. Eleven years later the queue is still alphabetical. Her model of the organisation is

accurate. The organisation's model of itself is the one that is wrong.

Now put her back in Chapter 1, thirty-seven years old, tool in hand, doing arithmetic about whether to raise it. I described that arithmetic as though she were performing it fresh — weighing IT, Legal, Quality Systems.

She wasn't. She ran that calculation at twenty-six. Chapter 1 is not deliberation. It is recall.

. . .

The cost of staying down

Hiding is expensive, and it is expensive continuously. A book that only tallied the costs of speaking up would make the submarine look far more comfortable than it is.

Conservation of resources theory holds that people work from a finite pool of psychological resources, and that maintaining effortful states draws that pool down. Concealment is such a state. Maria keeps a private account, strips data before it moves, declines to expense eighteen euros, avoids explaining how she caught the seal defect early, and tracks, in every conversation about her work, which version of the truth she is currently in. None of it is difficult. All of it is constant.

The emerging work on covert AI use at work finds exactly this: hiding your AI use draws down psychological resources in ways that show up as strain and depletion. It converges with a much older finding about concealment generally, which is that keeping something hidden at work costs the hider something real, every day, for as long as it stays hidden.

So the submarine has a running cost. It is paid in attention and energy by the person crewing it and it never stops.

Which raises the question this chapter exists to answer. If hiding is expensive, and the organisation genuinely wants the idea, why does the expense not eventually push people up?

Because the alternative is worse, and worse in a way that lasts.

. . .

What the research calls this

There is a large, old body of work on this, and it sits nowhere near the innovation literature. That is one reason the innovation literature has been so bad at explaining the fleet. It sits in organisational behaviour, under employee voice and silence.

Elizabeth Morrison and Frances Milliken's 2000 paper in the *Academy of Management Review* is the foundational statement. Their move was to argue that organisational silence is not an absence of something. It is a produced state – something organisations actively create through their structures and their managers' beliefs, and then fail to notice they have created.

Three years later Milliken, Morrison and Hewlin interviewed employees about why they stayed quiet. The large majority had at some point felt unable to raise a concern. Only about half felt generally comfortable speaking up at all.

The useful part is not the prevalence. It is the breakdown of reasons. The most common was fear of being labelled – a complainer, a troublemaker, someone difficult. Close behind, damaging a relationship. And in third place, at roughly a quarter of cases, something different in kind: the belief that speaking up would change nothing.

Hold those apart, because the rest of the chapter turns on the distinction.

Fear is the belief that speaking up will hurt you. **Futility** is the belief that speaking up will do nothing.

From outside they look identical, since both produce a person who says nothing. They are not remotely the same thing.

Fear is the loud one

Fear is what everyone means when they discuss this problem, and the literature has treated it thoroughly.

Amy Edmondson's work on psychological safety is the canonical treatment. Her 1999 study of hospital teams established the construct and showed that a team's shared belief about whether it is safe to take interpersonal risks predicts whether people speak up, and downstream, whether the team learns. The construct has since accumulated meta-analytic support linking it to performance.

Chapter 2 gave us the sharpest version of why the fear is rational: Allen and colleagues found employees who used shadow IT were marked down whether their tool succeeded or failed. The penalty attached to the circumvention, not the outcome.

Fear is real, documented, and has a known remedy. The remedy is safety — leaders who respond well to bad news, teams where the interpersonal risk of speaking is genuinely low. Building it is hard, and Edmondson has spent a career on how, but the mechanism is understood and the industry has absorbed it.

Futility is the quiet one, and it does the damage

Here is the problem. Maria was not afraid.

Look for the fear in her story. Her team leader liked the idea and passed it up. Nobody labelled her, no relationship was damaged, and nothing in that account suggests she was taking a risk. By any reasonable audit she was in a psychologically safe environment. She went silent anyway.

Safety answers *will this hurt me?* Maria's question was *will this do anything?* Safety does not answer the second question. It was never designed to.

Which is why well-intentioned interventions in this space fail while their sponsors remain genuinely puzzled. The most common is the open door — *my door is always open, bring me anything*. It is offered sincerely. It addresses fear. It does nothing whatever about futility, because the door was never the obstacle. What is on the other side of the door is the obstacle. An open door onto a room where nothing happens is not an improvement on a closed one. It is a more polite way of teaching the same lesson.

The same logic dismantles nearly every mechanism organisations reach for. Suggestion boxes. Innovation portals. Idea challenges. Skip-levels. Each is an invitation to speak, and each assumes people are silent because they have not been invited. Bernstein's factory workers, from Chapter 1, were concealing better methods a few metres from suggestion boxes that stood empty. They had not failed to notice the boxes.

It compounds

If futility were only a mood, this would be an argument about tone. What makes it structural is that repeated exposure changes the person.

Andrieu and colleagues exposed participants to repeated instances of voice that produced no result — non-instrumental voice, in the literature's phrase — and then measured not what those people said next, but how well they thought. Performance on a cognitive task fell. Motivation fell. Perceived control fell. What did not rise was any sense of being tired or distracted: the people affected did not report feeling worn down, they simply performed worse. Futility does not just close a mouth. It degrades the equipment.

That result sits on top of one of the most examined findings in twentieth-century psychology. Learned helplessness describes exactly this shape: repeated mismatch between action and outcome produces passivity, and the passivity generalises well beyond the situation that produced it. The organisational finding is new. The underlying mechanism is not.

Two further findings complete the picture.

Pinder and Harlos named the end state: acquiescent silence, distinguished from other kinds by the fact that nothing is being withheld. The person has given up hope of improvement and feels powerless to change things, so there is no suppressed idea to recover. Withheld silence still contains something. Acquiescent silence has stopped producing.

And Detert and Edmondson explain how the lesson persists without ever being re-taught. Their work on implicit voice theories found that people carry taken-for-granted beliefs about when speaking up is risky or pointless, and that these beliefs operate before any conversation begins. The person self-edits in advance. They do not weigh this idea against this manager, because the weighing already happened, generically, years ago.

Which is what Chapter 1 was. Maria did not evaluate her manager that Tuesday. She consulted a conclusion she reached at twenty-six.

. . .

The thing that should worry an executive

When a good idea is killed — or worse, allowed to stall — the organisation records the loss of one idea. Usually it records nothing at all.

What it lost is the person's future ideas. Not this one. All of them. The person has updated, the update is about the class of action rather than the instance, and it does not expire.

Do the accounting slowly. Maria at twenty-six would have produced, over a career, some number of improvements of the kind she is capable of. She is a good engineer, she notices things, and she had eleven more years at that plant ahead of her. The alphabetical queue cost the company one sort order. The lesson cost it everything downstream of a person who concluded, correctly, that offering things is not how this works.

And none of it is visible. That is a structural property, not a rhetorical flourish, and it is why the mechanism has survived sixty years of people writing about innovation. There is no record of an idea that was never offered. No system captures it, no metric moves, nobody files a report saying *I thought of something today and decided not to mention it*. The rejections you can count are the visible tip. The extinction is invisible by construction.

Engagement surveys will not find it either, because the extinguished employee is not disengaged. Maria is excellent at her job. She catches seal defects early. She would score fine.

I think this is among the largest costs a large organisation carries, and it appears on no ledger anywhere. I am aware of the bind in asserting that: I have just spent three pages arguing the cost cannot be measured, which means I cannot produce a number for it either. I am asserting it anyway, because the alternative — declining to name a cost because the organisation has no instrument for it — is precisely the failure this book is about. An unmeasured cost is not a small one. It is an unmeasured one.

Why fixing the gate does not fix this

Here is the property that makes learned extinction the answer to *why hasn't anyone fixed it*, and it is the hardest idea in the book.

Suppose you read Part II, accept every word, and go fix your gate. You separate two-way doors from one-way doors. You give increments a route that does not demand a business case. You defuse the territorial defence. You do all of it properly, and by Monday morning your organisation is genuinely better at receiving small ideas than it was on Friday.

Almost nothing happens.

The constraint was never only the gate. The constraint is a population that has already learned not to approach it, and that population does not re-learn on Monday because the gate changed. They will not test it. That is what the implicit-voice-theory finding means: the self-editing happens before the conversation, so the improved gate is never sampled. Your new process sits there, genuinely good, unvisited.

Engineers will recognise the shape. It is hysteresis — the state of a system depending not only on current conditions but on the path that got it there. Push the system into a new state and it does not return along the route it came. The damage persists after the cause is removed.

This is the best answer available to the question a sceptic should have been asking since Chapter 2. If all of this is so well documented, why has nobody fixed it? Partly because the mechanisms are individually reasonable, which was Part II. Mostly because the fix has a lag measured in years while the damage stays invisible, so nobody can tell whether their repair is working. An organisation that fixes its gate and sees no change within two quarters concludes the repair failed, reverts, and is wrong.

That is a hard thing to ask of an executive: do the expensive structural thing, receive no visible feedback for a long time, and trust that a mechanism you cannot measure is slowly reversing. I do not have a comfortable answer. I have a partial one in Chapter 12, and it is smaller than the problem.

Two mechanisms, one person

I have separated fear and futility cleanly, and Maria's story isolates futility on purpose, because a mechanism is easier to see when nothing else is running.

Real cases are usually both. Chapter 1 Maria is plainly afraid — of IT, of Legal, of an eighteen-month programme that would find her tool embarrassing. Chapter 9 Maria, eleven years earlier, was not afraid of anything. Same person, two mechanisms, and the second one arrived after the first had already taught her the ground rules.

The separation is diagnostic, not natural. It matters because the remedies do not substitute for one another. Safety without responsiveness produces a comfortable room where nothing happens, which is the open door. Responsiveness without safety produces a manager who acts decisively on ideas nobody dares bring. Most organisations attempt the first and call it culture work.

There is also a population this chapter does not describe. Some people surface repeatedly, get nothing, and surface again. They exist in every company and you can probably name one. Seniority explains some of it, since a principal engineer can absorb a stalled idea that would extinguish a twenty-six-year-old. Temperament explains more. What I will say is that an organisation relying on those people is not solving the problem, it is being subsidised by the small minority whose response to futility is unusual. That is not a strategy, and it is not fair to them.

But staying down does not work either

I have spent a chapter arguing that surfacing is dangerous and that the damage from a bad surfacing lasts a career. Stop there and I would have written an argument for permanent submergence. So let me close that door, because it is the other half of the resolution.

The submarine fails too, on a longer timescale, for different reasons.

Maria's tool exists in one place, which is Maria. It has no documentation, because documenting it means explaining it. It has no second user, because a second user is a witness. It runs on an account the company cannot see and does not know exists.

Nonaka's work on organisational knowledge creation supplies the vocabulary. Knowledge that stays tacit — held in a person, expressed in how they work rather than in any artefact — cannot be transferred, combined or built on. An organisation's capacity to create knowledge depends on tacit knowledge being made explicit and shared. Maria's tool is tacit by design, and permanently so.

So when she transfers next spring, three things happen at once. The tool leaves. The seal defects go back to being caught five weeks later. And nobody connects the second to the first, because nobody knew about the first, so the degradation reads as ordinary variation.

This is the documented failure mode for improvement work generally. The continuous-improvement literature is consistent that local gains decay unless they are institutionalised — sustained through documentation, standardisation and organisational support rather than through one person's commitment. There are widely quoted practitioner figures about how much decays and how fast. I am not repeating them,

because I cannot trace them to anything solid and this book does not launder folklore into fact. The direction holds regardless: unsurfaced improvement is not permanent improvement. It is a loan against one person's tenure.

Which gives the problem its real shape, and it is cruel.

Hiding costs the person continuously and loses the improvement eventually. Surfacing risks losing the improvement immediately and, done badly, costs the organisation every future improvement that person would have offered.

Both roads lose. The submarine is a bad permanent home and the surface is a dangerous place to be born.

Something has to give, and what gives is the assumption that this is one decision. It is two. The conditions an idea needs to survive its birth are not the conditions it needs to survive its adulthood, and the organisations that get this right stop asking people to choose between them.

That is Chapter 12. First we formalise the choice itself, and then we deal with the thing every executive wants to do the moment they finish this chapter, which is to go and find the fleet.

. . .

Where the evidence stops

The voice-and-silence research is decades deep and it is not in dispute. Morrison and Milliken established that silence is produced rather than absent. Pinder and Harlos identified the end state. Detert and Edmondson showed the self-editing runs before the conversation starts. Edmondson built the safety literature that explains the half of the problem everyone already knows about. Andrieu and colleagues took it a step further under controlled conditions, showing that repeated futility

degrades executive functioning itself, sitting on a learned-helplessness foundation as solid as anything in psychology.

Every one of those describes silence. None of them describes hidden innovation. Twenty-five years of work on why people do not raise concerns, and almost none of it pointed at why people build things and say nothing — which is a different behaviour growing from the same root.

The hysteresis argument runs past the evidence, and it is worth being exact about where. The literature establishes that implicit voice theories form and that they operate pre-emptively. It stops there. Nobody has tested whether those beliefs update when conditions improve, which is precisely the question that decides whether a repaired gate ever gets sampled.

So here is the study I would like someone to run. Take an organisation that has measurably improved how it handles small ideas, and count the lag before submission rates move. If they move inside a quarter, the persistence argument in this chapter is wrong and Part III weakens with it. My expectation is that the lag runs to years.

Maria illustrates the mechanism. The research carries it. Those are different jobs and I have not confused them.

. . . .

Surface Check

So — *should Maria raise her hand on Tuesday?*

Now we know *why* the question is hard for her specifically, and it is not courage. She ran this experiment before and the result was clean: nothing happened, slowly, and nobody was to blame. That is still the best model she has of how her organisation responds to being offered things. If she raises her hand and it goes the same way, she does not lose the tool — she confirms the lesson, and this time it holds for good.

If she stays down, the tool dies with her transfer next spring and nobody knows enough to miss it.

Both roads lose. Which means the question was wrong. Not “should she surface?” but “surface into what?” — and that is the only version with an answer.

INSTRUMENT

The Futility Audit

Harder than the earlier instruments, and the one I most want a manager to actually do.

Part one. The last five.

Write down the last five ideas anyone brought you. Not proposals that arrived through a formal process — ideas, including the ones raised in a corridor or at the end of a meeting when the agenda ran out.

For each, answer three questions:

1. What happened to it? “Taken forward” is not an answer. “Went to the March council and was funded” is, and so is “I meant to look at it.”

2. **Did the person find out what happened, and how?** Not whether an outcome exists — whether they were told. Silence from your side reads as futility from theirs, whatever was actually going on.
3. How long between them raising it and them hearing something?

If you cannot remember five, that is the finding. It means you are not being brought many, which is what this chapter is about.

Part two. The uncomfortable one.

For each of the five: **has that person brought you another idea since?**

This is the measurement. Everything above it is context. A manager who receives ideas from the same three people every quarter and never from anyone else does not have three creative employees. They have a department where everyone else has already run the experiment.

Part three. The distinction.

For anyone who has gone quiet, work out which mechanism you are looking at, because the remedies do not substitute:

- If they are *careful* around you — hedging, checking your mood, framing everything as a question — that is fear, and safety is the remedy.
- If they are entirely comfortable with you and simply do not bring things, that is futility, and safety will not touch it. Only demonstrated responsiveness touches futility, repeatedly, over long enough for someone to update a belief they have held for years.

The second is more common than you want it to be, and the fix is slower than any programme cycle you have ever run. That lag is why this problem is still here.

Next: Part IV. The dilemma, made rigorous — and why the curtain in Bernstein's factory matters more than anything else in this book.

Maria is still sitting at her desk. The meeting is Tuesday.

PART IV

Surface or Submerge

Everything up to here has been diagnosis, and diagnosis is the easy half.

Part IV is the decision, and it gets taken twice — once by the person with something to declare, and once by the person who would be declared to. Four chapters, and the camera moves between those two chairs rather than settling into either.

I will say now what the answer is not. It is not surface everything, and it is not leave them all down there. Both are available, both are wrong, and the first is wrong in a way that costs more than doing nothing at all, which is the least intuitive claim in this book and the one I have given a whole chapter to.

We start with the decision as the person underwater actually experiences it.

CHAPTER TEN

The Revelation Threshold

On the Friday, four days before the meeting, Maria showed it to somebody.

Not to anyone who would be in the room. To the man from logistics — the one who said the thing about pallet labelling at the first Improvement Time session four years ago, which she is still using, and whose name she had to think about for a second before she said it. They ended up at the same table in the canteen. He asked how the deviation reports were going, in the way people ask about work they do not do, and she heard herself say that she had a way of doing them now. He looked interested. She offered to show him.

They went back to her desk after lunch, and before she could open anything she had to say what it was.

She found she had no way of saying it. Nine months, and the thing had never had a name. In her own head it is *the report thing*, and on her laptop it is a folder called *dev* with a date after it, and neither of those survives being spoken out loud to a colleague standing behind your chair. So she said: it is a little tool I put together for the deviation reports. That is true, and it is the first sentence anybody in that building has ever heard about it, and she chose it in about a second and a half.

Then she showed him the wrong thing.

Not by accident and not because she meant to hide anything. She opened the output rather than the tool, which is the summary it produces on a Friday morning, and she talked him through that, and twice she said the word *rough* before he had said anything at all. When he asked how it decided what to put at the top, she said it looks for repeats. Which is not what it does. It clusters by failure mode across suppliers, and the reason that is the whole point is that the standard report is organised by supplier, and explaining why that ordering is wrong would have taken ten minutes and would have amounted to saying that the way this plant has read its deviation data for as long as anyone can remember is the wrong way round. That is a criticism of a system, and a system belongs to people, and two of them sit on her floor.

So she showed him a faster version of the report he already knew.

He was impressed by it, which was worse than if he had not been, because what impressed him was the wrong thing — the speed. He said it must save her hours. She said it did. He asked whether it was on the network and she said no, it is just on here, and he nodded in a way that closed the subject, and she registered that he had understood something she had not said.

Then he said the dangerous thing, and he said it kindly.

He said she should show it to the man with *excellence* in his job title. Not as a warning and not as a hint. As the obvious next step, offered by somebody being helpful, because that is what you do with a thing that works — you take it to the person whose job is things that work. He is right by every ordinary standard, and Maria has spent eight chapters arriving at the reason he is wrong, and she said something about maybe, after the review, and he let it go.

On the way out he asked whether anyone else knew.

She said no. He said she should tell them. And then the conversation moved to something about a trailer that had been sitting in the yard for three days, and they walked back through the corridor with the poster in it, and nothing further was said about any of this.

What changed that afternoon was not the tool. The tool is exactly what it was at half past eleven. What changed is that there is now a person in the building who could say one sentence about it in a room she is not in — not out of malice, not even out of carelessness, just conversationally, in a meeting where somebody asks whether anyone is doing anything about deviation reporting. Chapter 9 was about what a bad surfacing does to the person for the rest of their career. This is the term before that one. Before the sanction can be applied there has to be something to apply it to, and until Friday lunchtime there wasn't.

That evening she did something she has not done in nine months, which is write the decision down. Two columns on the back of a printed agenda.

She could not fill in the left one, because Chapter 6 already told you what is in that box and it is not knowable yet. The right one she could fill in, and the first thing she wrote at the top of it was not *they say no*.

It was: *they take it*.

There is a formal decision rule for exactly this, published last year in an innovation journal, with terms and a threshold and a small amount of algebra. Maria arrived at it in a canteen without writing anything down, which is roughly the relationship this whole book has to the literature.

. . .

The rule, already written down

The cleanest statement of the surface-or-submerge decision is not mine, and it is barely a year old. In 2025 Julian Waters-Lynch, Darcy Allen, Jason Potts and Chris Berg published a paper on what they call shadow user innovation — covert use of consumer AI tools inside firms, positioned deliberately at the junction of the user-innovation literature, the shadow-IT literature, and the bootlegging work this book opened with. It is the most on-point external source I have found for the whole subject, and it does something none of the older literatures do. It writes the employee's choice as a rule.

Their model has two triggers and one threshold. Covert use starts where there is acute pain in a work process and where the fix is concealable. It ends — that is, the person discloses — when the expected gain from disclosing exceeds the expected sanction risk. Below that line they stay covert. They call the line a revelation threshold.

Read as a description of a person it is almost insultingly obvious, and that is a recommendation rather than a complaint. Good formalisations of familiar behaviour usually look obvious. What the formalisation buys you is that it separates the decision into terms, and terms can be examined one at a time, and the moment you examine them one at a time against the previous nine chapters something becomes visible that is not visible in the prose version.

The rule is correct. Every input to it is wrong in the same direction.

. . .

Four terms, all bent the same way

Take them in order, and notice as we go that none of what follows requires anybody in Maria's plant to be unreasonable.

The gain, at the moment of asking. The expected benefit of disclosure has to be estimated by the person disclosing, on the day they disclose. Chapter 6 spent itself on why that estimate is systematically low: the value of a good increment is produced by running it, in the place it runs, over a period, and the largest part of it is usually something nobody predicted. Maria in month zero could not have priced the seal defect because in month zero the tool did not catch seal defects. So the gain term is not merely uncertain. It is biased downward by construction, and biased hardest for the increments most worth having, which is the paradox from Chapter 6 restated as an input to somebody's arithmetic.

The gain, conditional on a yes. Here is where the equation quietly assumes something false. It reads as though a favourable response converts the expected gain into the actual gain — as though the yes delivers the thing. Chapter 8 established that it does not. The yes arrives from a person who is good at his job, wants the plant to do well, and is structurally unable to receive a small working thing except by giving it a sponsor, a scope and a name. The smallest version of Maria's tool that the organisation is capable of adopting is larger than the tool. So the gain has to be discounted not by the probability of rejection but by the probability that acceptance is of a different object.

The sanction, which does not discount for being right. Chapter 2 gave the sharpest version of this and it is worth repeating, because it is the finding that most people refuse to believe until they have watched it. Allen and colleagues found that employees who used unsanctioned tools were marked down whether the tool succeeded or failed. The penalty attached to

the circumvention. A rational person facing that evidence cannot reduce their sanction estimate by building something excellent, because excellence is not the variable being punished, which means the one lever the engineer actually controls is disconnected from the term it ought to move.

The sanction, which does not expire. And this is Chapter 9 arriving as a coefficient. A bad surfacing is not a one-off cost paid on a Tuesday. It is an update to a belief the person will carry into every subsequent decision of this kind, for years, without ever re-testing it, because the self-editing runs before the conversation starts. Maria at twenty-six paid this once. What she has been carrying since is not a memory of a stalled proposal but a permanently raised sanction term, and it is raised for every idea she will ever have, not for the one she lost.

Four terms. Two of them understate the gain, two of them overstate or extend the risk, and all four errors push the same way — down, toward concealment.

Which gives the fleet its real explanation, and it is not the one most executives arrive with. The submarine is not crewed by people who are frightened, or disengaged, or insufficiently included in the vision. It is crewed by people doing correct arithmetic on inputs the organisation has bent.

That is a considerably worse problem than a morale problem, because you cannot fix it by being nicer to anybody. But it is at least a problem with a shape. If the terms are bent by the machinery, then the machinery can in principle be unbent, and Chapter 12 is an attempt at that.

There is one more thing wrong with the equation, though, and it is not a bent input. It is a missing term, and it is the reason this chapter exists rather than being a paragraph in the last one.

. . .

The curtain

I have been holding this since Chapter 1, and it is the strongest piece of evidence in the book.

I should say plainly what I mean by that, because this book has been careful about the difference between an analogy and a result. Almost everything in Parts II and III arrives here from somewhere else. The capital research is about plants and equipment. The attention research is about subsidiaries. The selection research is largely about chief executives of listed companies, and I flagged the size of that leap where I made it. Tyre and von Hippel and the NUMMI study are the two bricks that sit on an actual assembly line, and they are case studies.

What follows is a controlled result, obtained on a production line, in a working factory, by a researcher who went in looking for the opposite of what he found. There is nothing else like it in this book and I am not going to pretend the rest of my evidence is of the same kind.

In 2012 Ethan Bernstein published a paper in *Administrative Science Quarterly* based on work inside the world's second-largest mobile phone factory. He went in to study transparency, which was then and is now one of the few management ideas nobody argues with – that the more of the work managers can see, the better the work will be. Visibility improves quality. Visibility catches error. Visibility spreads good practice from the line that has it to the line that does not.

What he found on those lines was that the operators had worked out better ways of doing the job than the official ones. Faster, and in some cases more reliable. Improvements of exactly the kind the previous nine chapters have been about, invented in place, by the people with their hands on the work, without any of it having been asked for.

They hid them.

When a manager came down the line, the improvised methods went away and the operators returned to the procedure in the manual, and they went back to the better version once the coast was clear. This was not occasional and it was not a few troublemakers. It was how the floor worked. And the suggestion boxes standing on every line, which existed precisely so that this would not have to happen, stayed empty — which Chapter 9 has already explained, and which at the time must have looked like a workforce with nothing to say.

Then Bernstein hung a curtain.

Not a metaphor. Fabric, around parts of the production floor, creating zones the operators could work inside without being observed from outside — with the output still measured, and the line still accountable for what it produced. He reduced what managers could see, on purpose, in a factory, and left everything else in place.

Performance improved, and it stayed improved.

I want to be exact about the direction of that, because it is the part people reverse when they retell it. The gain did not come from the workers being happier, or trusted, or motivated by a gesture. It came from the improvements being allowed to exist. The better methods already existed before the curtain; what the curtain removed was the reason to put them away, and the line got the benefit of work its own people had already done and had been hiding for as long as anyone had been watching.

Bernstein named the finding himself, and his name for it is the one this book will use. He called it the transparency paradox.

His prescription is narrower than the story usually told about it, and the narrowness matters enough that I will not smooth it over. He is not arguing for factories nobody looks at. He argues for bounded privacy, with specific boundaries — privacy around teams rather than individuals, a separation between feedback and evaluation, a separation between the right to decide and the

right to improve, and protected periods for experimentation. Observation is not the enemy in his account. Observation wired directly to judgement is. That distinction is doing quiet work here and it will do loud work in Chapter 12.

. . .

The principle, and where you have already met it

Now go back and look at what happened in the canteen.

Maria was not afraid of the man from logistics. He has no authority over her, he had already been helpful to her once, and nothing about that afternoon carried any risk that would show up in the equation above. And within ninety seconds of him standing behind her chair she had opened the output instead of the tool, described the clustering as *it looks for repeats*, apologised for the state of it twice, and presented the whole thing as a faster version of the report the plant already has. She showed him the by-the-book version. She had a curtain and she took it down herself, and what came out from behind it was the manual.

That is Bernstein's operators, in an office, with a laptop, nine months in.

And it is the third time this has happened in this book. The first was in Chapter 6, at the form. *Expected benefit (quantified)* did not measure Maria's tool badly. It measured it accurately, in month zero, when it clustered by supplier and saved her an afternoon — and month zero is the least the thing will ever be worth. The looking did not misread the value. The looking chose the moment, and the moment set the value.

The second was Chapter 7, at programme scale. Nobody killed the protected time by judging it and finding it wanting. Somebody sensibly asked how many people were using it, and somebody sensibly asked what had come of it, and each request

was reasonable, and together they rebuilt the gate inside the sanctuary. The audit did not discover that the programme had stopped being a nursery. The audit is what stopped it being one.

Three altitudes — an idea, a programme, a person — and one mechanism running at all three.

Observation does not reveal the state of a hidden improvement. It sets it. And it sets it downward.

That is the second organising claim in this book, alongside the three-currency convergence from Chapter 5, and I will call it the observation principle, on the understanding that the name is a container for Bernstein's finding rather than a rival to it. He established the mechanism at the scale where it can be measured. The claim that the same mechanism is operating at the other two scales is mine, it is an argument rather than a result, and the caveats section deals with how far it runs past him.

But notice what it buys, because it is more than tidiness. Everything in Part II mis-valued the increment downward, and never once upward, and that asymmetry has been sitting there unexplained since Chapter 4. Here is the explanation. Looking early collapses a thing to whatever is legible on the day of looking, and legibility correlates with size, with precedent, and with distance from the actual work — which is Chapter 6's closing line about a system that quietly optimises for forecastability and calls it rigour. The direction of the error is not an accident of any particular gate. It follows from the act of measuring an embryo.

. . .

Four results, and none of them is the same study

I have just built an organising claim on one factory in one country, and the caveats section will say so again in harder

terms. Before it does, it is worth laying out what else is standing near it, because the situation is better than one study and worse than a settled literature, and the difference between those two matters.

Four results, none of which was designed to test any of the others.

A factory. Bernstein's lines, where privacy raised output and the effect held for months. A field experiment with a control, in the setting where the work is most visible and most standardised.

A laboratory. Amabile's 1990 crossing of observation and evaluation, twenty-two years earlier, where the expectation of judgement — not the watching — reduced the creativity of what people produced. Different method, different decade, different participants, no organisation anywhere in it.

A workgroup. Brown and Baer's territorial marking, from Chapter 4, where the act of one party staking a claim measurably suppressed the creativity of everyone else in the room. Nobody rejected anything. The environment changed and the output changed with it.

A decision meeting. Mueller's bias against creativity, also from Chapter 4, where the evaluator's own uncertainty degraded the reception of exactly the ideas most worth having, and did it below the level at which the evaluator could feel it happening.

What they share is not a hypothesis. It is a direction. In all four the conditions of being seen or judged altered what was produced or offered, rather than revealing what was already there, and in all four the alteration ran the same way — toward the safe, the conventional, the defensible. Four methods, four scales, one direction.

What that does not do is prove the observation principle. These are not four tests of one claim, and I am not going to present them as if they were; three of the four were run for

other reasons and would not have been designed this way. Convergence of this kind is weaker than replication and stronger than a single site. What it buys is a specific and limited thing: if Bernstein's factory turned out tomorrow to be a fluke, the direction would still have three independent results behind it, and the argument would be wounded rather than dead.

. . . .

What that does to the equation

The revelation threshold treats the improvement as a constant. There is a thing, it has a value, and the decision is whether to reveal it — with the same object sitting on both branches, unaffected by which branch is taken.

The observation principle says that is exactly what it is not. Disclosure is not a report about the tool. It is an event in the tool's life, and it is an event that changes what the tool is: what gets shown, how it gets described, what it will be permitted to become afterwards. Maria's tool on the concealed branch clusters by failure mode and catches things five weeks early. Maria's tool on the disclosed branch, described in twenty minutes to somebody standing behind her chair, is a faster report — and if the room adopts anything on Tuesday, that is what it adopts, because that is what was in the room.

Two branches, two different objects. So the comparison the equation asks for cannot be made in the form it asks for it.

Which does not make the threshold useless. It makes it a rule about timing rather than a rule about worth, and that is the more useful thing to have. The question stops being *is this good enough to show* and becomes *has this become what it is going to be*, and the second question, unlike the first, has an answer that changes with the calendar.

So the question has a date on it

Chapter 6 explained why delayed assessment works, and it is worth pulling the practical consequence out of it, because nine Surface Checks have now said *not yet* and I have never said what the *yet* is waiting for.

Criscuolo, Salter and Ter Wal found that hidden projects came out better, and the mechanism was time — the embryo gets to become what it is going to be before anyone judges it. Nothing in that finding says hidden forever. It says hidden until, and the paper is not in a position to tell you when the until arrives, because it measured outcomes rather than moments.

The practitioner literature is more direct about it and less rigorous, which is the usual trade. Miller and Wedell-Wedellsborg's *Innovation as Usual*, and I will flag that as a trade book rather than a study, argues for stealth in the early stages on the grounds that a premature no is close to irreversible — once a proposal has been declined, everything subsequently learned about it gets read through the decline, and the people who declined it have a position to defend. Their advice is to stay under the radar until the thing can survive contact.

Put that against the observation principle and the shape of an answer appears, though not a comfortable one. If observation sets the value, and it sets it at whatever is legible on the day, then the only thing under the discloser's control is which day. Not whether to be seen. When. And the honest test for the day is whether the thing has accumulated something legible that is also true — a rate you can defend from your own records, a second person who has used it and would say so, an event it has already survived. Maria has the first of those, has had it since about month four, and has neither of the others.

The Surface Checks have been saying *not yet* for nine chapters. What they have been waiting for is a version of the tool that can be seen without becoming worse, and nothing in her plant is building one, and she cannot build it alone, because two of the three things on that list require somebody else.

. . .

The decision in its real shape

Part IV opened by promising to make the dilemma rigorous, and the rigour has turned out to cut in an awkward direction. The decision rule is real and it is not the book's. Waters-Lynch and colleagues wrote it down properly: disclose when the expected gain beats the sanction risk, otherwise stay under. Everything this book has assembled goes into that rule as inputs, and every one of the inputs is bent the same way. The gain is estimated on the day the value has not yet been produced. The gain on a yes is discounted by the fact that a yes converts the object into a larger one. The sanction does not fall when the work is good, because the penalty attaches to the circumvention. And the sanction, once paid, is not paid once.

Then the principle takes the last piece of solid ground away, which is the assumption that there is one object being decided about. There is not. There is a thing that exists in the dark and a different, poorer, more legible thing that exists in the light, and the act of moving between them is what makes the second one out of the first.

So Maria's arithmetic is sound and her conclusion is correct and both of those facts are a catastrophe for the plant, because the plant needs the tool and there is no move available to her that delivers it. That is not a failure of nerve on her part and it is not a failure of goodwill on anybody else's. It is a system in which the only way to see the thing is to spoil it.

Which is where the executive enters, holding the obvious solution, and it is the worst idea in this book.

If people conceal, look harder. If the concealment is invisible, get better instruments — and there has never been an instrument for this like the one every large company has just finished buying. That instinct is not stupid, it follows directly from everything above, and it is precisely inverted. The observation principle is now on the table as a principle rather than as a story about a curtain, and Chapter 11 is what it predicts.

. . . .

Where the evidence stops

The load in this chapter is carried unevenly, and the unevenness is worth mapping before somebody maps it for me.

Bernstein is one factory. It is a Chinese contract-manufacturing plant, on assembly lines, with a defined standard method that operators were deviating from — which is the cleanest possible setting for this finding and not obviously the same setting as a quality engineer with a laptop and no standard method to deviate from. The step from an operator hiding a hand movement to Maria hiding a clustering rule is a step. I think it is a short one, because the thing being concealed in both cases is a local improvement to a prescribed process and the trigger in both cases is being watched by someone who could act on what they see. But it is one site, one industry, one country, and I am leaning on it harder than any other single source in the book. The convergence section above is what stops that being fatal, and it is worth being clear that it does not stop it being a weakness.

His prescription is narrower than my use of it. Bernstein argues for bounded privacy with four specific boundaries, not

for reduced visibility as a general good, and executives have entirely legitimate reasons to see things — capital, risk, compliance, coordination. Where this chapter says *observation sets the outcome*, his finding says something more precise: observation wired to evaluation and control sets the outcome. I have tried to keep that distinction visible, and Chapter 12 depends on it completely, because a resolution built on the loose version would be a resolution nobody could operate.

The revelation threshold is a model, not a measurement. The Waters-Lynch paper is a theory piece. It proposes a decision rule and derives testable propositions from it; it does not report a test. Nobody has measured a revelation threshold in a working organisation, established what actually sits in the terms, or checked whether people behave as the rule says they do. I have used it as scaffolding for exactly that reason — it organises the argument, and I have been careful not to write as though it settles anything.

Miller and Wedell-Wedellsborg is a trade book, said once already and repeated because it is the kind of source that gets promoted to a study by a helpful reader. The stealth-innovation advice is practitioner reasoning drawn from consulting work. The confirmation-bias mechanism behind it is well established; the specific advice about timing is not tested.

And the three-scale claim is mine. Bernstein established the person scale. Criscuolo and colleagues supply the programme scale indirectly, in the finding that the same bootlegging behaviour associated with *less* innovation during periods of process formalisation — same behaviour, observed surroundings, inverted result — which is suggestive and is not the same as having been tested as a mechanism. The idea scale, Chapter 6's form, has no controlled result behind it at all. It has an argument and a good one. Nobody has assembled these as one principle because nobody has had reason to, and if the

assembly is wrong, it is wrong in the way arguments are wrong rather than in the way findings are wrong, which is to say you can dispute it without new data.

Here is the study that would settle it. Take a large organisation with a population of tools people have built for themselves — every big company now has hundreds, and most of them do not know it. For each one, record two dates: when it was built, and when somebody with authority first saw it. Then look at all of them a year later and ask which still exist, which are used by more than one person, and what they are doing. If the tools that were shown in month one are in the same state a year on as the tools that were shown in month nine, then the observation principle is wrong outside a production line, Bernstein's result does not travel, and this chapter is a nice story about a curtain.

My expectation is that early disclosure predicts one of two endings — the tool stops, or the tool becomes a project and then stops — and that late disclosure predicts survival.

The obvious objection to that study is the one I would raise myself, and it is fatal to the cheap version of it. People disclose early because the thing is already good, so any raw comparison measures the tools rather than the timing. To get anything out of it the timing has to be assigned rather than chosen, which means a real experiment inside a real company, which means somebody senior agreeing to keep half of it deliberately unobserved for a year. That is a hard thing to ask for. It is also exactly what Bernstein got somebody to agree to in 2012, which is the other reason that study is worth more than everything else I have.

. . .

Surface Check

So — should Maria raise her hand on Tuesday?

No, and for the first time the reason has nothing to do with the room. She has already run the experiment, in a canteen, with a friendly colleague and no authority present, and within ninety seconds she had reduced her own tool to a faster version of the report the plant already has. The room would not spoil the thing on Tuesday. She would, in the act of describing it, because that is what being looked at does.

Which finally makes the question answerable. It is not whether to surface. It is what would have to exist on the other side for surfacing not to cost her the thing — and nobody in that building is building it, and she cannot build it by herself.

INSTRUMENT

The Revelation-Threshold Worksheet

The first instrument in this book aimed at the person holding something rather than the person receiving it. Twenty minutes, on paper, and not on a company device.

Part one. Name it and date it.

Write down what you have built, in one sentence, as if to somebody who does not do your job. Notice how long that takes. Most people find they have never said it out loud, and the first sentence you produce under pressure is the one that will be in the room for the rest of the thing's life.

Then two dates: when you started using it, and when it started being worth something. They will not be the same date and the gap between them is the argument of Chapter 6 in your own handwriting.

Part two. The gain, as a rate.

Not a lump sum. How many times a year does it run, what does one run save or catch, and what is that over three years? Rates are defensible from records; lump sums are estimates and get treated as such.

Then write the thing it has done that you did not build it to do. If there isn't one yet, that is a finding about timing, not about quality.

Part three. The yes you should be afraid of.

Name the person most likely to take it up. Then answer one question about them, factually and without generosity in either direction: **what is the smallest version of this that they are able to receive?**

Not the smallest they would like. The smallest their position allows — the smallest thing that can have their name on it and still be worth their attention. If that version is bigger than what you have built, then acceptance and destruction are the same event, and no amount of preparing your case changes it.

Part four. The sanction, and the part of it that does not expire.

Write the worst realistic outcome. Most people find it is smaller than the dread suggests.

Then the question that actually matters, and it is the one from the last chapter: *what would that outcome do to the next one?* Not to this tool. To the one you have not had yet. If your honest answer is that you would stop offering things, then the cost of a bad surfacing is not the tool, and you should price it accordingly rather than pretending the risk is bounded.

Part five. A date, not a verdict.

The point of all of the above is not to reach a yes or a no. It is to work out what would have to be true for the answer to change.

Three things make a hidden improvement survivable in daylight, and you can only be sure of the first:

1. A rate you can defend from your own records.
2. A second person who has used it and would say so.
3. An event it has already survived — a bad week, a data problem, a case where it was wrong and you caught it.

Write down which of the three you have, and the date you will ask yourself this again. Then keep going.

Part six. For the person on the other side of the table.

If you are the one being surfaced to, this worksheet has a short version, and every question in it is factual.

What is the shortest route in your area from a working thing to a supported thing that does not pass through a business case? If you cannot describe it in one sentence, there isn't one.

Has anybody used that route in the last twelve months? Name them.

And the last one, which is the useful one: think of the most recent time somebody showed you something before it was finished. What did you do in the following week — and would they say the same?

If you cannot remember an unfinished thing being shown to you, that is not evidence that nobody is building anything.

Next: Chapter 11 — the instinct every executive has at this point, and what the evidence says it costs.

Maria is still sitting at her desk. The meeting is Tuesday.

Stop Hunting the Fleet

On the Monday, the man with *excellence* in his title opened a report he had not asked for, and found the thing this book has spent ten chapters describing.

Some of it, anyway. He had asked one question, six weeks earlier, on a call with three other plants, and the question was whether anybody actually knew what tools people were using. It was not an accusation. It was the sort of thing you say when a slide about AI readiness has been on the screen for four minutes and you cannot think of anything else to contribute. Somebody from the group function wrote it down, and six weeks later a discovery report arrived in his inbox with the plant's name at the top of it and a covering note saying this was now available monthly.

Sixty-one people. Out of a device estate of about four hundred.

Sixty-one people on his site had used a generative AI tool through a company machine in the previous ninety days, and none of it had gone through anything. Not a licence, not a request, not a ticket. He read the number twice, and what he felt was not anger, which is worth being precise about, because the story that gets told afterwards always installs anger at this point

and there wasn't any. What he felt was the specific relief of a man who has been asking why nothing moves faster for two years and has just been handed a reason.

Then he read the names, because they were in the appendix and nobody who receives a list of names does not read it. He knew most of them. Three were in his own function, one of whom he had recommended for a development programme in March. There was a maintenance planner he had never spoken to. There was a scheduler who had put a query into a chat tool about how to phrase a supplier email, which the report had classified alongside everything else, because the report did not know the difference and was never built to.

Maria was not on it.

She is not on it for a reason nobody designed and nobody would defend if it were put to them. The report finds sign-ins to a list of consumer AI products, because that is what everybody pictures when they picture this problem — a person with a chat window open. Maria's tool does not have a chat window. It is a script that makes a call to an interface, on an account in her own name, paid for with eighteen euros a month she has never expensed, and from the outside the traffic is indistinguishable from any other request leaving any other laptop in that building. The most developed case of the thing the report exists to find sits four hundred metres from the man reading it and produces no row.

Nor do the others in that building doing versions of the same thing, one of whom appears in the sixty-one for something unrelated and trivial.

What the report can see is the maintenance planner, who used the machine on his desk because it did not occur to him that he was doing anything.

That is the whole of the argument in this chapter and I have put it in the cold open rather than saving it, because everything

that follows is an attempt to establish that the shape of that list is not an accident of one plant's rollout. It is what the instrument does. It is what any instrument of that kind does, and the better the instrument gets, the more reliably it does it.

He spent about forty minutes on it and then did something entirely defensible. He wrote a line for Tuesday's deck — *establish visibility of unmanaged AI usage* — and under it two sub-points, one about a licensed platform and one about a communication to line management. He was not planning to name anybody. He had no intention of punishing anyone and he would have said so, plainly, if asked. He thought he had found a gap and he was proposing to close it, which is the job he has and the job he is good at.

The next report is due in a month and the number will be lower.

He will read the fall as the communication working. And this book's position, which is the position the rest of the chapter has to earn, is that the fall will mean nothing of the kind — that the number went down because the work went somewhere his instrument cannot follow, and that he will have bought a picture that improves as the thing it measures gets worse.

. . .

From the other side of the desk

That is the first time in eleven chapters we have sat in that chair, and it is deliberate, because the argument in this chapter is not survivable from Maria's. From her side it writes itself and writes itself badly. Surveillance arrives, people hide, management is the problem, and the reader who already agreed nods and the reader who did not stops. I am not interested in that version. The man in the cold open has not done anything wrong by any standard I can articulate, he is better at his job than most people

reading this, and he is about to make the single most expensive mistake available to him — not through carelessness but by following the evidence he has to the conclusion it supports.

So I want to give his instinct its strongest form before I take it apart, because a version of it that is easy to refute is not the version anybody actually holds. It goes like this. Chapters 9 and 10 have just established that people are concealing useful work, that the concealment is rational, and that it is invisible by construction. Nothing about that is disputed. And the first thing any competent manager does with an invisible problem is try to make it visible, because you cannot govern, resource, protect or even count a thing you cannot see. Nobody should have to defend that. It is the definition of the job.

There is more, and it is the part the critics of monitoring tend to skip. Some of the looking is not optional. If plant data is being pasted into a consumer service, there is a data-protection question with a legal answer, and that answer does not soften because the person pasting had good intentions. If a process is under a quality standard with an audit trail attached, an unrecorded tool in the middle of it is a real exposure, not a theoretical one. Licensing, retention, access, the ordinary machinery of a regulated manufacturer — all of that gives an executive legitimate reasons to want to know what is running in his building, and none of them are cover stories.

I will come back to that class of reason in Chapter 12, because separating it from the rest is most of the work of building anything that functions. Here I want to grant it and set it aside, since it is not what the report on Monday morning was for. That report was not commissioned to answer a data-protection question. It arrived because somebody asked what people were using, and it will be spent on making sure fewer of them use it.

And the instinct now has something it never had before, which is why this chapter belongs to the present decade rather than to any of the previous six.

. . . .

What the principle predicts

Knowledge work was never protected from monitoring by policy. It was protected by cost.

You could watch a warehouse floor because the work was already legible — pick rates, scan times, distances, all of it generated as a by-product of the work itself, in a form a computer could count without anybody having to interpret anything. Call centres for the same reason. Delivery, obviously. What made a quality engineer unwatchable was not that anyone thought she should be left alone. It was that reading her afternoon would have required a person to sit and read it, and no organisation has ever had that person spare.

That constraint is gone, and it went quietly, in about two years, as a side effect of buying something else. The same capability that lets a company summarise a thousand support tickets lets it summarise a thousand employee sessions, and the second use requires no new purchase and, in most estates, no new permission. There is now, for the first time, an instrument that can watch judgement work at the resolution at which a warehouse has always been watched.

Which is the moment to say what the last chapter established, because the prediction falls out of it rather than being brought in from outside. Observation does not reveal the state of a hidden improvement. It sets it, and it sets it downward. Bernstein's operators put the better method away when a manager came down the line; Maria opened the output instead of the tool within ninety seconds of a friendly colleague standing

behind her chair; a form asking for expected benefit fixed the value of a nine-month tool at what it was worth in month zero; an audit converted a sanctuary into the gate it had been built as an alternative to. Four instances, three scales, one mechanism.

An always-on instrument pointed at everybody is that mechanism running continuously, at every scale at once, with nobody in a position to arrange a moment when they are not being looked at.

So the prediction is specific and it is not about morale. First, the improvements will not appear, because what a watched person produces is the version that survives being watched, and Chapter 10 spent itself on why that version is the poorer one. Second, the performance gain that justifies the purchase will not arrive, because the gain was always supposed to come from seeing the better methods and the better methods are the first thing to go. Third, the concealment does not stop. It moves.

The research that follows is consistent with all three, and I want to be exact about the epistemic order here, because it matters more than the citations do. The meta-analyses confirm this. They did not produce it. If the monitoring literature did not exist, the prediction would still follow from Bernstein plus the argument of the last chapter, and it would still be worth acting on. What the literature does is stop it being one factory.

. . .

And it is not about delivery riders

Here is the objection I would raise, and it is the right one. Almost everything vivid in the monitoring literature comes from platform and gig work. Delivery riders gaming a dispatch algorithm. Warehouse ethnographies of work games and workarounds. The algorithmic-management studies, which are the closest published analogue to the argument in this chapter,

are largely drawn from platform populations and largely from Chinese samples. A quality engineer with a laptop, a permanent contract and fourteen years of tenure is not a rider, her situation is not that situation, and any argument that leans on riders to predict her behaviour deserves to be asked why.

The answer is that the evidence on office work exists, it is meta-analytic, and it runs the objection backwards. Electronic performance monitoring has been studied on administrative and office populations for about thirty years, and it has been pooled. One meta-analysis covering roughly seventy samples and more than two hundred separate effects found what you would broadly expect on the human side: monitored employees are slightly less satisfied and slightly more stressed, and there is a small increase in counterproductive behaviour. The effects are small and I am not going to inflate them.

The finding that matters is the one missing from that list.

Across those samples the relationship between being monitored and performing better was indistinguishable from zero. A larger pooling — ninety-four studies, more than twenty-three thousand participants — reaches the same place: no overall link between monitoring and better performance. That is not a finding about how monitoring makes people feel. It is a finding about the thing the instrument was bought for.

So the objection — *this is all about riders, my people are different* — is true, and it points the wrong way, though it takes one more step to see why. Those pooled samples are mostly clerical and administrative populations. That is the work monitoring was built for: countable, comparable, close in shape to a rider's shift. And the performance gain the instrument is bought for does not appear even there. Maria's work is not that work, and the part of it this book is about — the part that produced a tool nobody asked for — has its own evidence, which is older, more direct, and not a null at all.

That evidence is from 1990, and it is the thing in this chapter I find most persuasive, because nobody involved was thinking about software. Teresa Amabile and colleagues crossed two variables people usually run together: being watched, and expecting to be judged. Being watched, on its own, did comparatively little. Expecting the work to be evaluated did a great deal, consistently, across a verbal task and an artistic one.

The work came back less creative. Not less productive. Less creative — more conventional, safer, closer to the method you would defend if asked. And what damage surveillance did appears to have run through the same channel: it was not the watching, it was what the watching implied was coming. That is the same shape as Bernstein's line and the same shape as Maria in the canteen, arrived at twenty-two years earlier, in a laboratory, with no factory anywhere in sight.

Which is the answer to the rider, and it is worth stating in the order the evidence actually runs. On the work monitoring was designed for, thirty years of pooling has not found the performance gain. On the discretionary work the executive believes is different, the one controlled result we have says the output degrades — not from being seen, but from expecting to be judged. The objection is right that Maria is not a rider. It is wrong about which way that cuts.

The mechanism has a name and a long literature. Psychological reactance describes what people do when a freedom they had is threatened: they move to restore it, and the restoration is frequently indirect. Monitoring threatens autonomy directly, autonomy is one of the conditions under which people do discretionary work at all, and the restoring move available to a salaried employee is almost never confrontation. It is the quiet relocation of the thing that matters to somewhere the threat does not reach.

Two more pieces close the gap to AI specifically. The first is direct. Liu, Yuan and Jiang published work in 2025, combining a situational experiment with a multi-wave field study, finding that algorithmic management increases knowledge hiding — that people managed more by algorithm withhold more from the people around them. The mediator is psychological ownership: the more the system treats your knowledge as the organisation's to extract, the more firmly you hold it as yours. It is the transparency paradox with software in the manager's chair.

The second is about when the trigger fires, and it is the distinction this chapter and the next both hang on. A synthesis of the monitoring research published in *Harvard Business Review* in 2024 — a synthesis in a practitioner outlet, not a study, and I will say that once — pulls out the variable that separates monitoring that helps from monitoring that costs. It is not how much you watch. It is what the watching is wired to. Where monitoring data feeds development, the harms are muted. Where it feeds evaluation and control — performance conversations, reviews, decisions about people — counterproductive behaviour rises.

Which is Bernstein again, from a different direction, and it is worth saying flatly because the popular version of his finding gets it wrong in a way that makes this chapter sound like something it is not. He did not argue for factories nobody looks at. He argued for bounded privacy with specific boundaries, and his own account of what went wrong on those lines is not that managers looked. It is that looking and judging were the same act.

The report that arrived on Monday morning is wired to judgement. It was wired that way in the ninety minutes it took to write the covering note, by nobody in particular, because that is the default wiring and choosing anything else requires somebody to have read a paper.

The fleet relocates

Now the part with the least evidence behind it and the most at stake, and I would rather flag that at the top than have it noticed later.

The sixty-one will fall. On this I am confident, and so is every practitioner I have ever heard discuss it, and the confidence is not the argument. The argument is about where the difference goes.

Nothing in the report changes whether the deviation reports are still painful on a Friday afternoon. Waters-Lynch and colleagues put two triggers on the front of their model: covert use begins where there is acute pain in a work process and where the fix is concealable. A monitoring rollout does not touch the first trigger at all. What it does is change the price of the second — it raises the return on concealability, sharply, for everybody who was previously not bothering to conceal anything.

That is a derivation from their model rather than a result they report, and it predicts a channel shift rather than a stop. The tool leaves the managed laptop and goes to the phone. The company account becomes a personal account. The work that used to happen at a desk on Thursday afternoon happens on a sofa on Thursday evening, on a home network, in a browser nobody has ever inventoried, and the person doing it now has a reason to keep it that way that they did not have in September.

They are not being defiant. The pain is still there and the route to fixing it just acquired a cost.

There is one place where somebody has already run this experiment on a different channel and the result is on the public record, in dollars.

Financial firms are required to preserve business communications, and for most of them the control that delivers it is a policy naming the permitted channels and banning the rest. The policy existed and was breached at scale.

Staff conducted business on personal devices and consumer messaging applications, none of it captured, and regulators have since charged more than a hundred firms — a hundred and twenty-five million dollars against JPMorgan in December 2021, a combined one point eight billion across sixteen firms in September 2022 with something like six hundred and seventy-six thousand communications never archived, and further waves through 2025 taking the running total past three billion.

The part that matters here is that most of the penalised firms already had the ban.

The prohibition was the control, and the finding against them was a finding about an absence the prohibition had helped create. The work went where it could not be seen, for the most ordinary reason available, which is that clients wanted to be texted. Then generative AI arrived in February 2023 and the same industry reached first for the same instrument, before spending the following three years building the sanctioned internal tools it could have built at the start.

I am borrowing that evidence and not the authority behind it. I do not work in financial services, an enforcement record is not a controlled study, and off-channel messaging is a records problem rather than an improvement one — nobody's WhatsApp thread was going to catch a seal defect five weeks early. What it does establish is that the relocation I am predicting is not hypothetical behaviour invented to make an argument work. It has happened, at scale, in an industry with more compliance resource than any factory has ever had, against a clear prohibition, with a statutory penalty attached to being caught.

We know the destination exists because it is already the main road. The MIT NANDA work in 2025 found roughly forty per cent of companies holding official large-language-model subscriptions and roughly ninety per cent of employees using personal AI tools for work — and I have already been careful with that comparison once, in Chapter 1, and the caution stands. What it establishes is not that shadow beats official. It establishes that the personal channel is not a bolt-hole somebody would have to find. It is where most of this was happening before anybody bought an instrument, and a rollout does not have to create it. It only has to make it more attractive.

So count what the plant has after the number falls. It has less visibility than before, not more, because the fraction that was running on managed devices — the maintenance planner, the scheduler, the man in Maria's function who was going to mention it eventually — has moved off them. It has no more control than before, and arguably less, because the one route by which any of this could have become official ran through people being relaxed about being seen. And it has acquired a data-governance position that is worse in every direction: plant data on consumer accounts, outside retention policy, outside the data-processing agreement, outside anything the security function could enumerate if regulators asked, and with no incident-response path because there is no incident anyone can detect.

The instrument was bought, in part, to reduce exposure. Its first-order effect is to increase it, and to make the increase invisible.

That is the flip I would put in front of a board, and there is nothing moral in it. It is the argument the security function would make about any control that pushes traffic off the estate, and it happens to be the strongest thing this chapter has to say to the person actually holding the budget. You have not sunk the

submarine. You have pushed it into international waters, where you have no jurisdiction, no telemetry and no standing to ask what it is carrying.

. . .

The number goes down and it means the opposite

Which produces a failure mode that deserves naming on its own, because it is the reason this mistake does not get caught. The report measures activity on monitored channels. Concealment moves activity off monitored channels. So the metric improves, the underlying state worsens, and the two are not merely coincidental — the same event causes both. The better the compliance number looks, the more of the estate has gone dark.

Every organisation on earth has a version of this and it is normally taught as a curiosity about targets. Here it is load-bearing, because it means the mistake is self-funding. Six months in, the number is down, the intervention appears to have worked, and the reasonable next step is more of it. Nothing in the feedback loop is capable of reporting the cost. A programme that was wrong in month one accumulates evidence of its own success at exactly the rate it does damage.

And the damage has the property Chapter 9 established, which is that it does not reverse when the policy does. Suppose the man with *excellence* in his title reads this chapter in a year, believes it, and switches the whole thing off. The sixty-one do not come back. They have learned something about what happens to work that can be seen, they learned it from an event they can name, and the self-editing runs before any future conversation the way it always has. He will have paid the full price and then paid again for the repair, and his submission rates will not move for years, and nothing in his reporting will tell him why.

It finds the people you least want to find

There is one more property of the instrument and it is the one I have not seen written down anywhere, so treat it as mine.

Monitoring is a selection instrument, and like every selection instrument it has a bias, and the bias is knowable in advance. It cannot select for the most hidden work, because that is what hidden means. What it selects for is carelessness and convenience and the assumption that one is not doing anything that needs concealing, which is to say it selects, with some precision, for the people who were closest to surfacing anyway.

Look again at the list from Monday morning. The maintenance planner is on it because he used the machine in front of him. The scheduler is on it because it did not occur to her that asking how to phrase an email was a governance event. The man in Maria's function is on it because he had been intending to raise it at some point and saw no reason to be careful in the meantime. Every one of them is a person for whom the revelation threshold from Chapter 10 was already close to being crossed — low estimated sanction, low perceived need to hide, one conversation away from the thing becoming official.

Maria is not on it, and the proximate reason is an accident of what the report was scoped to look for. The reason underneath it is not an accident. She crossed that calculation eleven years ago and has been living inside the answer ever since — a private account, the data stripped before it moves, a receipt she has never filed.

So the instrument reaches into the population and removes, first and most reliably, the people who were about to give you what you wanted. It converts the least-hidden into the properly hidden. And it does so by teaching them the lesson in the most

efficient possible way — not through a rejection, which at least requires them to have offered something, but through an unsolicited discovery that arrives with a policy attached and makes clear they were being watched the entire time.

The people it does not reach are the ones who were never going to tell you anyway, and after the rollout there are more of them, and every one of them has been handed a fresh reason.

. . .

The hypothesis, stated plainly

This is the book's central claim and the point at which the foreword's promise comes due, so I will put it in one place without hedging it and then say exactly which parts of it are new.

Pointing AI downward, as an instrument for finding hidden work, destroys more value than the hidden work was costing you. It does so through four steps, none of which requires anybody to behave badly. Being watched sets the state of a hidden improvement, downward, which is the observation principle and is documented at the scale where it can be measured. Monitoring returns no performance gain at all, across ninety-four studies, which is meta-analytic and is the single most-ignored result in the whole field — and where the work is discretionary, expecting evaluation has been shown to degrade it. Concealment relocates rather than stops, which follows from a published model of the disclosure decision and is the least-tested link in the chain. And the instrument selects preferentially against the people who were nearest to disclosing, which compounds through the extinction mechanism from Chapter 9 and does not reverse when the policy is withdrawn.

Take those together and the sentence that has been sitting under this book since the foreword becomes sayable. The only way to see the fleet is to stop hunting it.

That is not a plea for trust and I have no interest in arguing it on those grounds. It is a claim about what an instrument can and cannot return, made about a specific instrument, in a specific decade, and it is falsifiable, and the section after next says how.

Two things it is not, because both misreadings are available and both would be worse than the mistake this chapter is about. It is not an argument that executives should stop looking. Bernstein did not stop looking; the output was still measured and the line was still accountable for what it produced, and if that had not been true his finding would be worthless. What he changed was what the looking was wired to. That distinction is the entire difference between a resolution somebody could operate and a resolution nobody could, and the next chapter is built on it.

And it is not an argument that the submarine is fine where it is. Chapter 9 closed that door: the hidden tool dies with the person who built it, and Maria transfers next spring, and nobody in that plant will connect the seal defects going back to being caught five weeks late with a folder called *dev* that left on a laptop. The fleet has to come up. Everything in this chapter is about what happens when you try to make it come up by force.

There is a way to do it that does not use force, and it is not a culture programme, and it has an unglamorous, structural, faintly boring shape that I think is the only thing on offer that actually works. Which is Chapter 12.

. . .

Where the evidence stops

More of this chapter is argument than in any other, and the load is distributed unevenly, so here is the map.

The effect sizes are small. The monitoring meta-analyses report modest reductions in satisfaction, modest increases in stress, and a small rise in counterproductive behaviour. I have deliberately not built on those, and anybody who wants to defend monitoring on the grounds that its harms are minor is arguing against a case I have not made.

Half the case is a null, and nulls are weaker ground than results. The claim that monitoring returns no performance gain is a claim about an absence, pooled across studies of varying quality, on populations that are mostly clerical and administrative rather than professional. An absence of measured effect is not proof of no effect, and heterogeneous pooling can hide real gains in particular settings. What I will defend is narrower and still sufficient: nobody has demonstrated the performance benefit that monitoring is routinely purchased to deliver, and thirty years is a long time to have not demonstrated it.

And the other half is a laboratory. The affirmative claim — that expecting evaluation degrades discretionary work — rests on Amabile, and Amabile ran short tasks with student participants rather than a nine-month side project in a plant. The step from an artistic task under an evaluation instruction to a quality engineer hiding a clustering rule is a step I am making. It is the same step Bernstein's factory lets me make at the person scale, from the other direction, which is why I am willing to make it. It has not been tested where it matters.

The algorithmic-management evidence is adjacent, not identical. Liu and colleagues studied knowledge hiding — withholding knowledge from coworkers — which is a cousin of

concealing a project from management rather than the same behaviour. The samples are largely platform and Chinese, the effect is indirect, and it is moderated: it bites hardest where organisational identification is low. That moderation is honest and it cuts both ways. It weakens any claim that the effect is universal, and it says the effect is strongest exactly where the fleet is largest.

The relocation claim has no direct measurement at all. This is the weakest link and it is not close. I have a published model with concealability as a trigger, survey evidence that the personal channel is already the majority route, a well-established mechanism in reactance, an enforcement record from another industry showing the behaviour occurring against an explicit ban, and a great deal of practitioner agreement. The enforcement record is the strongest of those and it is still the wrong shape: it establishes that people relocate work when a channel is prohibited, not that they relocate it when it is monitored, and those are different stimuli. What I do not have is a single study that watched an organisation introduce monitoring and then measured whether the work moved to personal devices, because as far as I can establish nobody has run one. The chapter's most consequential claim is a prediction.

Nobody has studied the actual case. There is no research on knowledge workers concealing AI-assisted process improvements from AI-assisted monitoring. Every component is evidenced and the assembly is not, and the assembly is where the argument lives. This is the newest ground in the book and I would rather say so than let the citation density imply otherwise.

Here is the study that would settle it. Take an organisation about to introduce monitoring of AI tool use, and measure two things before and after — not on-channel usage, which will fall and tells you nothing, but the proportion of people who report

using personal devices or personal accounts for work tasks, and the rate at which employee-built tools are voluntarily shown to a manager. If on-channel usage falls, personal-channel use does not rise, and voluntary disclosure holds steady, the relocation claim is wrong, the compliance number means what it appears to mean, and this chapter is a long argument against something that works.

The problem with that study is the same problem as everything else in this book, and I find it clarifying rather than discouraging. You would be asking people to self-report, accurately, on a behaviour that the intervention has just made punishable. Run it in-house and the answers are worthless. It needs a third party, genuine anonymity, and a sponsor willing to fund a result they will not enjoy.

A phenomenon that resists being measured by the act of measuring it is not an excuse for a weak evidence base. It is the thesis, arriving in the methodology section.

. . .

Surface Check

So — should Maria raise her hand on Tuesday?

The question has changed shape overnight and not in her favour. Until Monday morning the worst realistic outcome of surfacing was Chapter 8's outcome: her tool becomes a workstream and stops being her tool. There is now a second one. A report exists that she is not on, a monthly cadence has been announced, and anything she says on Tuesday about a personal account and eighteen euros a month lands in a room where the framing was set on Monday by a document about unmanaged usage.

She will not be describing an improvement. She will be describing a finding.

And the timing is not something she chose, or could have anticipated, or can now do anything about. Nine months of correct arithmetic, and the term that decides it moved four days before the meeting, for reasons that have nothing to do with her, in an inbox she has no access to.

INSTRUMENT

The Surveillance Backfire Check

For the person considering the instrument, or already holding it. Every question below is factual and answerable from records or from a five-minute conversation, and none of them asks you to rate your own culture, because you cannot and neither can anybody else.

Part one. What decision changes?

Write down the decision that will be made differently because of this data. Not the goal — the decision. Who makes it, when, and what are the options in front of them.

If you cannot name one, you have not bought an instrument. You have bought a camera, and the rest of this check will not help you, because a camera with no attached decision still changes behaviour and returns nothing in exchange.

Part two. Trace the wiring.

Four factual questions, and the answers are in your own systems.

1. Who can see individual-level output from this data? List them by role.
2. Does anybody on that list write, contribute to, or sign off a performance review for anybody who appears in it?
3. Is the data retained past the point where the decision in Part One gets made? For how long, and who decided that?
4. If somebody appears in it tomorrow, what is the first thing that happens to them, in the order it actually happens?

If the first thing that happens is a conversation about policy, the instrument is evaluative. That remains true if the conversation is friendly, if you did not intend it that way, and if the word *support* appears in the email.

Part three. The relocation ceiling.

Which of the following can your instrument see? Answer for the tool as configured, not the tool as sold.

- A personal phone on the guest network.
- A personal account in a browser tab on a managed laptop.
- The same laptop, at home, on a home network.
- A photograph of a screen.
- A prompt typed on a phone and the answer typed back in by hand.

Count the noes. That is the ceiling on what you can ever learn, and it does not move with a better product. Now ask which of those five gets more attractive the week your rollout is announced.

Part four. Pre-register the fall.

Before it goes live, write down and date what you will conclude if the number drops by half in six months.

Then, on the same page, write down everything else that could produce an identical drop. There are at least three and you will find them in twenty minutes.

Six months from now the number will have moved and you will be told a story about it. This page is the only thing that will stop the story being the one you already wanted.

Part five. The question that costs something.

Before the instrument goes live, find three people who are not on any list — not your direct reports, not the people who bring you things already — and ask each of them to show you something they have built or worked out for themselves.

Not in a meeting. Not with a form. Ask, and then wait, and do not fill the silence.

Whatever comes back is the baseline you will never be able to reconstruct afterwards. If nothing comes back at all, note the date, because that is not an argument for a better instrument. It is the finding the instrument was going to cost you the ability to see.

Next: Chapter 12 — the shape of the thing that does work, which is smaller and duller than the problem and is the only version I can defend.

Maria is still sitting at her desk. The meeting is Tuesday.

The Dual Operating System

On the Monday evening, the man with *excellence* in his title went down two floors to find somebody, and he had not worked out on the stairs what he was going to say when he got there.

The report had arrived that morning and he had given it forty minutes and written the line for Tuesday's deck before eleven. What he had not finished with was the appendix. Three of the sixty-one names sat in his own function, and one of them belonged to a process engineer he had put forward for a development programme in March — a nomination he had written himself, three paragraphs, and he could remember using the word *initiative* in the second one because he had thought at the time that it was a slightly lazy word and had left it in anyway.

Seeing that name under a heading about unmanaged usage did not make him angry either. What it made him feel was that he had recommended somebody on the strength of an impression, and that the impression had turned out to be correct in a way he had known nothing about, which is a stranger thing to sit with than being wrong.

The engineer was still at his desk with his coat already over the back of the chair. He is thirty-four. He came off the floor, which matters to all of what follows — six years as a setter on

the presses before he did the qualification, and people who have done that job know, in a way that cannot be arrived at from an office, exactly which two steps of a changeover are the ones where it goes wrong.

He asked him what he had been using it for. He asked it badly, because he had not found a way of putting the question that did not sound like the first line of a disciplinary conversation. The engineer heard that, picked his coat back up, and said it was easier to show him.

So they walked out to the line at twenty past five, which is how the head of a function comes to be standing at a seal press in the second half of a shift, being shown something nobody had ever intended him to see.

The board is on a stand at the end of the machine, at the height of somebody standing at it, and it is the kind of board that exists to carry the shift's numbers. There were no numbers on it. What was on it, in marker, in two colours, was the changeover procedure — the steps in the order they actually happen rather than the order the 2011 document lists them in, in German down the left-hand side, and beside it, line for line, the same steps in Turkish.

Two stations along, on another board, the same thing in Czech.

At the bottom of the German column, under a line he had ruled himself with the edge of a clipboard, were two items that appear in no controlled document anywhere in that plant: the two points in the changeover where it goes wrong, which he knows because for six years he was the person it went wrong for.

In the top right corner of the board there was a post-it.
NICHT WISCHEN.

Underneath it, a second one, older, the adhesive gone brown at the edges, with the same two words on it in a different hand.

The executive asked about the post-its, which is the question anybody asks, and got the answer that the board had been wiped once in February by somebody doing a housekeeping sweep at the weekend who could not possibly have known what it was, and that putting it back had taken most of a Saturday shift, and that the second post-it had been written a fortnight later by a team leader on nights who did not trust the first one to be enough.

Then he asked the obvious thing, which is how a process engineer from the office comes to be writing Turkish on a board. The engineer said he does not speak Turkish. Nobody in the office does. He puts the controlled instruction into a chat tool on the machine at his desk, gets back a version in plain German with the steps reordered, asks for it again in Turkish and in Czech, and copies it onto the board by hand, which takes about twenty minutes and which he does on a Tuesday, usually because somebody on shift has asked him the same question for the third time and he would rather answer it once where they can all see it.

Eleven stations since April.

Asked whether any of it is written down anywhere that is not a whiteboard, he said there is a file. Somebody started it four years ago on a shared drive, and the night shift types things into it when there is time, and things get carried from one shift to the next in it, and he is not certain who owns it or whether anybody does. The two failure points from the bottom of the board are in that file, somewhere around row sixty, in German only.

He said all of this without preparing an account of it and without appearing to think that one was needed, which is the part the executive kept returning to afterwards. Asked why it had never come up, he said he had not thought of it as something that came up. Then he said, without being asked and

without any change in his voice, that he did not know what would happen if somebody from the auditor's side stopped at the board.

He was right to wonder, and the executive knew precisely what would happen, and it is worth being exact about it, because nobody in this story is behaving badly. A document that controls a safety-relevant changeover exists in one approved version, carrying a revision number and a signature, and that is not bureaucracy for its own sake — it is the mechanism by which a plant can say afterwards what a person on shift was told to do on a given day.

What is on that board is a second version of a controlled instruction. It is unapproved, it is machine-translated, and there is nobody in the quality department who can read the Turkish column to check that it says what the German says. The finding writes itself. The corrective action is a cloth, and the auditor who writes it up is doing the job the standard exists to make them do.

And the board is better than the document. It is in the language the crew reads, in the order the work happens, and it carries two pieces of information about how that press fails that exist nowhere in the plant's controlled documentation and are otherwise held by one man and an unowned spreadsheet. Both of those things are true at the same time and neither cancels the other, and the executive stood in front of a whiteboard with two post-its on it and held both.

He went through what he could do about it, and there turned out to be four things. He could escalate it, in which case the board is wiped inside a fortnight, correctly. He could ask for the work instructions to be reissued properly in four languages, which is not a request but a project, with a scope and an owner and a budget line, and Chapter 8 has already explained what happens to the size of a thing that needs those. He could say

nothing, which had stopped being available to him about ninety seconds earlier. Or he could leave the line that was already sitting in tomorrow's deck.

Four routes, and every one of them ends with a clean whiteboard.

He did not think about Maria once that evening, and there is no version of the evening in which he would have. She is not in his function, she is not on the report, and in fourteen years nothing about her has ever arrived on his desk as a question — and a man who has just been handed a list of names does not go looking for the people who are not on it. The instrument missed her because of what it had been scoped to see. He missed her because she has never given him a reason to look, which is the same absence arriving twice by two entirely different routes.

The line stays in the deck. He thought about taking it out, on the walk to the car, and could not think what he would put in its place.

. . .

Two conditions, not one decision

Chapter 9 ended on a sentence that has been carrying a debt ever since, and this is where it gets paid. Both roads lose. Hiding costs the person continuously and loses the improvement eventually; surfacing risks losing the improvement immediately and, done badly, costs the organisation everything that person would have offered afterwards. I said then that what gives is the assumption that this is one decision.

It is two, and they are not merely different. They are opposed.

An improvement needs one set of conditions to survive being born and a different set to survive being twenty. At birth it needs privacy, no forecast, no owner other than the person building it, and time — the delayed assessment from Chapter 6, which is the

only thing that lets a tool that clusters by supplier in week one become a tool that clusters by failure mode in week six. Every one of those conditions is a condition of being unobserved, which is the observation principle from the other side. What protects the embryo is that nobody is looking at it yet.

Adulthood needs the reverse of all of it. Documentation, because Chapter 9's Nonaka argument says knowledge that stays in one head cannot be built on. A second user, because one user is a dependency. A name, an owner who is not necessarily the builder, a place it lives that is not a laptop, and somebody accountable when it is wrong on a Thursday. All of those are conditions of being seen.

So the resolution is a sentence about sequence rather than about culture, and it is the smallest thing this book has to offer.

Hide to survive birth, surface to survive adulthood. The submarine is a phase rather than a home, and almost every organisation runs one of the two states and applies it to everything.

There are exactly two ways to get this wrong and both are common enough to have a smell. The first is to apply the adulthood conditions at birth, which is Part II of this book in one line. The form in Chapter 6 asked for expected benefit in month zero. The gate in Chapter 4 asked for a case before there was anything to make a case about. Both were reasonable, both were the wrong test at the wrong time, and both produced the same result, which is that the thing either dies or goes under.

The second is to apply the birth conditions forever, which is the failure this book is least often accused of enabling and should be. Give everyone a capable tool, say nothing else, and you have built an efficient factory for producing submarines.

They work. They are not documented, they have no second user, and they leave the building in a laptop bag when their author transfers, which is next spring in Maria's case and some

other spring in everyone else's. The continuous-improvement literature has been describing this failure in manufacturing for twenty years — local gains that regress when the conditions that enabled them turn out to have been one person. There are widely quoted figures about how fast that happens and I am not repeating them, for the reason given in Chapter 2, but the direction is not in dispute in any study I can find.

Which means democratisation is half an answer. It is the birth half, and it is the half that can be announced, and most companies stop there because stopping there produces a press release and the other half does not.

The other half is duller and it has a shape that the industry survey work keeps arriving at from a different direction. Organisations that get anywhere with this tend to formalise what their people were already doing rather than start from a platform decision, and the deployments that stick often began with somebody private who had worked it out and was then made a champion of it. Channel rather than ban. Which is the same conclusion Chapter 11 reached from the other end, since banning does not stop the work — it moves it somewhere you cannot see, and then you have lost the route as well as the visibility.

. . . .

What big programmes are for

I want to concede something properly here rather than in a subordinate clause, because a reader who has got this far has every right to think this book is an argument against large projects, and it is not. Big programmes are the correct instrument for a class of problem that increments cannot touch. You cannot crowdsource a retooling. A new line, a platform migration, a change to the physical layout of a plant, a system

that has to be replaced entirely or not at all — these have genuine capital indivisibility, they have coordination costs that only a programme structure can carry, and Chapter 5's arithmetic about lumpy investment is not a critique of them but a description of why they are correct. An organisation that tried to reach those outcomes by letting four hundred people improve their own work would reach none of them.

The two are for different physics. And each has a failure mode that the other one does not.

Big programmes fail by escalation and by theatre. They become too large to cancel, they acquire a roadmap that ships slides rather than product, and the eighteen-month programme in Chapter 1 is still eighteen months from delivering when Maria builds her tool on a Sunday. Increments fail the opposite way. They work, they are never written down, and they evaporate — which is Chapter 7's Improvement Time and the whole of the paragraph above.

So the target of this book has never been either category. It is the organisation that owns only one of them and forces everything through it. A company with only big programmes starves its increments, because an increment has no structural weight and therefore never attracts any attention, and Chapter 5 explained why that is a property of finite attention rather than a failure of anybody's judgement. A company with only democratised increments watches its gains decay and never lands the bets that need mass. Both are one-system organisations. They fail differently and they fail for the same reason.

Which raises the obvious objection, and it is the one I would raise, because this book has spent an entire chapter demolishing exactly this kind of proposal. Improvement Time was a second system. So was twenty per cent time. Both were announced by somebody senior, both worked, and both thinned out into the

past tense without anybody deciding anything, for the reason Chapter 7 gave — a cost line with no matching benefit line survives only as long as somebody with standing keeps defending it, which is to say about as long as that person's tenure. Chapter 8 then made it worse, by showing that any structure large enough to be visible acquires a floor beneath which nothing is worth its attention. A second operating system that is a programme, a centre or a slot in the calendar is a thing this book has already buried twice.

The version I am proposing survives that test by being smaller than the things that failed it. It has no address, no head count, no director, no budget of its own and nothing to put in a planning round, because none of what follows is a programme. It is wiring. Who may see individual-level output from a monitoring tool, what that seeing is attached to, whether a person may change their own work without permission, how much money sits with the person rather than with the approver — every one of those is a standing arrangement rather than an activity, and a standing arrangement has no annual defender because it makes no annual request. Nobody has to protect it in February. There is nothing to withdraw.

It is exposed to a different failure mode from the one that killed Improvement Time, which is a smaller claim than saying it is easy, and it is the one I can defend.

. . .

Who pays

The cost objection arrives at about this point in every conversation I have had about this material, and it arrives in a specific and reasonable form. Not everybody should be on the frontier model. You cannot hand ten thousand people an

unmetered account against the most expensive tier and expect the bill to be defensible.

That is true, and I want to give it away completely before doing anything else with it. Tokens are not free, the frontier tier is the expensive one by a wide margin, and a policy of putting everybody on the largest available model is a policy with a real number attached. *Leapfrog* argued at some length that the frontier is rarely the right tool for the routine task, that most of what a working engineer needs runs on something smaller, cheaper or open, and that the instinct to reach for the biggest model is usually an instinct about reassurance rather than about capability. Nothing here contradicts that. This book inherits it.

But watch what the cost argument actually produces when it is made in a room.

It almost never produces a cheap default that everybody may use freely. What it produces is a central team, an approval route, a quota, and a business case per use case — which is the capital gate from Part II wearing a different jacket. One-way-door caution applied to a two-way-door expense. And the moment that machinery exists, the scale threshold from Chapter 4 reappears with it, because only a use large enough to be worth an approval cycle can survive an approval cycle, and every improvement of the kind this book is about is smaller than that by construction.

Maria's whole tool runs on about eighteen euros a month. The meeting required to approve eighteen euros a month costs more than the eighteen euros in its first sitting, and that is before anybody writes the form. I am offering that as an illustration rather than as a figure — I have not costed it and neither has anyone else — but the shape of it is not controversial to anybody who has sat in the meeting.

The flip is the part worth having, and it runs against the intuition. Central control is what produces runaway bills, not

what prevents them. A person spending from a pool they do not feel spends without noticing, and a platform team asked to serve everybody optimises for coverage rather than for thrift, because coverage is what it is judged on.

Push the capability and a small real budget out to the person doing the work, and the person living inside that budget becomes efficient in ways no policy could specify — smaller model by default, cache the thing that gets asked twice, batch the overnight job, stop when the answer is good enough. That is the learning curve from Chapter 6 applied to spend, and it works for the same reason it works everywhere else, which is that the knowledge is where the work is.

So the shape is a cheap or open default that nobody has to ask for, a small owned allowance that the person controls and can exhaust, and a clear route to the expensive tier for the task that warrants it. This is stricter than central approval rather than looser. It just puts the discipline in the hands of the only person in the building who knows whether this particular call needs the big model.

. . . .

What keeps it honest

Four governors, and the chapter does not work without them, because a resolution with no limits is a slogan.

Not everything needs to surface. Some tools are one person's helper and should be left where they are. If somebody has worked out a way of drafting their own emails faster, that is not a governance event, it does not belong on a register, and hauling it into daylight produces nothing except a person who will not tell you about the next thing. Surfacing is a door offered, not a net thrown, and the test is whether anybody other than the builder would be worse off if it vanished tomorrow. Where the

answer is no, leave it alone. Chapter 11's report classified a scheduler asking how to phrase a supplier email alongside everything else, because it could not tell the difference and was never built to, and a human system that cannot tell the difference either is not an improvement on it.

Safety is not cost. Chapter 11 granted a class of legitimate reasons to look — data protection, audit trails under a quality standard, licensing, retention — and deferred them here, and separating them from cost-theatre is most of the work of building anything that functions.

There is a test that does most of it. A real constraint names a category of *data*: this may not leave the estate, that may not enter any model at all, cheap or expensive. A cost constraint names a category of *spend*: who may use which tier, and who signs. The first survives being made free — if the tokens cost nothing tomorrow, the data-protection limit is unchanged and every word of it still applies. The second evaporates. Run the objections in your own organisation through that and you will find they sort cleanly, and you will also find that the ones that evaporate are usually the ones being argued hardest.

And there is a hard line. In safety-critical and regulated work the gate is load-bearing and surfacing is not optional. The man with the whiteboard is inside that line, and the answer is not that he should be left alone, because he should not. An unverifiable translation of a safety-relevant instruction is a real finding, and the auditor who writes it up is doing their job.

But notice what the two-system question does to that case rather than what the one-system answer does. Under one system the only available move is removal, and removal is what happens, and the plant loses the only version of that instruction the crew can read. Under two, there is a second question available — how long does it take this organisation to turn a better page into the controlled page — and the answer to that

question is a fact about the plant rather than about him. If it is a week, there is a second system. If it is eleven months and a project number, there is one, and the whiteboard was the workaround for it.

And futility eats all of it. Chapter 9 established that the belief that speaking up changes nothing does more damage than fear does, lasts longer, and is harder to reverse. A protected route that produces no visible response is not a neutral thing that failed to help. It is a demonstration, delivered to exactly the people who were willing to try, that trying does not work — and it teaches the lesson more efficiently than never having offered the route at all, because these are the people who took it up.

Improvement Time did not fail because two hours were insufficient. It failed because the third session was cancelled and there was never a fourth, and everyone in the room learned something that Friday morning that no subsequent programme has been able to unteach.

. . .

Bernstein's four boundaries, made operable

Everything above is the shape. This is the part somebody can act on, and I take it from the same paper that carries the strongest evidence in this book, because Bernstein did not stop at the curtain. He set out what the boundaries around a zone of privacy have to be, and the reason his prescription is narrower than the story people tell about it is that he was describing a control system rather than a mood. There are four of them, and each translates almost directly onto the case this book is about.

Privacy around teams, not around individuals. The curtain went around a stretch of line, not around a person. Applied here, the immediate consequence is about the report from Chapter 11: a discovery tool that produces a row per named employee is

doing the thing Bernstein's factory stopped doing, and one that reports at the level of a function or a site answers the governance question without producing a list of names. The information an executive actually needs — is plant data leaving the estate, is there exposure under the standard, how much of this is happening — survives aggregation completely. The appendix is what does not, and the appendix is the part that was read.

Feedback separated from evaluation. This is the wiring question and it has a factual answer in every organisation. Does anybody who sees individual-level output from this data write, contribute to, or sign off a performance review for anybody who appears in it? Where the answer is yes, the instrument is evaluative regardless of what the covering note says, and Chapter 11's meta-analytic finding applies: control-use is where the counterproductive behaviour comes from. Where the answer is no, and where the separation is structural rather than a promise about intent, the same data can be developmental and the harms are muted. Nothing about the sensitivity of the data changes between those two arrangements. Only the wiring does.

Decision rights separated from improvement rights. This is the one that matters most and the one almost nobody has, and it is a single sentence: you do not need permission to change how you do your own work; you need permission to change how anybody else does theirs. Maria has never had that distinction available to her. Her tool changes only her own afternoon and it has sat under nine months of the same calculation as a proposal to change the plant's reporting standard, because her organisation makes no distinction between the two and processes both through the same gate.

Write the distinction down, for one department, and most of the fleet in that department has no reason to be submerged the following week. The improvement stays local, the person keeps

the right to make it, and the moment it starts changing somebody else's work it becomes a decision and enters the machinery — which is the correct place for it to enter, because at that point it is one.

Protected periods for experimentation. The trap here is the one Chapter 7 walked into, and the difference is worth being exact about. Improvement Time was two hours in a calendar, which is a cost line, which is why an escalation on the third Friday was always going to win.

A protected period attached to a piece of work rather than to a slot in the week has no cost line to defend: the thing is exempt from being assessed until an agreed point — a date, a number of runs, a first external user — and until then nobody asks what it is worth. That is Chapter 6's delayed assessment written as a rule rather than as an accident, and it costs nothing, and there is nothing in it for a planning round to remove.

Put those four together and you have a system that is smaller than the problem it addresses and duller than any of the programmes that have failed at it, which is the strongest thing I can say for it.

The improvement is born unobserved and stays that way until it has become what it is going to be. The person building it is not deciding whether to break a rule, because there is no rule against improving your own work. When it starts to matter to somebody else there is a route with a name, and the route is short, and somebody has been down it in the last twelve months. And the monitoring that a regulated manufacturer legitimately needs runs at a level that answers the governance question and produces no list of names, so nobody learns the lesson that Chapter 9 spent itself on.

None of that requires anybody in the building to be braver than they are.

The route, described

Everything above this point is wiring, and wiring has a defect I should name before somebody else does. Its virtue is the one I have already claimed. A standing arrangement makes no annual request and therefore has no annual defender, and that is why I think it survives in a building where Improvement Time did not. But none of it is a thing a person does on a particular afternoon. Somebody with a budget and an org chart can read this chapter, agree with every line of it, and find that nothing whatever has changed by Friday — because agreement is not an action, and everything above describes a state rather than a move.

So here are three moves, with parameters in them.

Each is small enough that one manager can grant it without escalating. None has an address, a head count, a director or a budget line, and none of them goes into a planning round, so the survival argument from the last section applies to these as it applies to the wiring. And every number below is a starting value rather than a finding. They are the least-evidenced material in this book, the next section says so again in stronger terms, and they belong in something that can be corrected by the people who run them, because a book cannot be updated and a practice can.

The Protected Window

The one to build first, and the one that goes closest to the mechanism.

Criscuolo and colleagues found that bootleg effort raises innovative performance, and their explanation was delayed assessment — not avoided assessment, delayed. An idea in week one cannot survive a gate review, and the reason is not that it is

bad. It is that it is young, and there is no path from the first version to the approvable version that runs through a meeting.

Which means the thing people are hiding to obtain was never permission to build; it is permission not to be judged yet. No organisation I know of offers that, so they take it by concealment, and the window is the legitimate version of what Maria has been stealing for nine months.

A manager grants an idea a fixed period — six weeks is my working number — with a hard cap on spend, five hundred euros or thereabouts, and one condition, which is that closing it costs nothing except the time already spent. In exchange the idea is exempt from assessment for the duration. No business case, no gate review, no steering committee, no status update, no alignment check against whatever large programme is currently running in the same territory.

What it requires at the end is one conversation. The person says what they learned, and the answer *it does not work* is a complete and acceptable one, which is the part most organisations will find hardest to actually mean.

The mechanism is doing something narrow and I want to be exact about it. It gives the person the single thing concealment was buying them, at a price the organisation can afford to pay openly, and it converts hidden work into visible work without applying the judgement that caused the hiding in the first place. Nothing else in this chapter does that in one move.

The open questions are real and I would rather set them down than pretend the design is finished. Six weeks may be wrong in either direction — long enough to reach a defensible version, short enough that a manager can grant it without asking anybody, and those two constraints may not overlap in every organisation. The end of a successful window is where I think this is most likely to fail, because the idea now has to enter the normal world and that transition is exactly the machinery the

window was built to defer. Renewal may quietly reintroduce what the window was avoiding. And the largest one: nothing stops this becoming a programme with an application form, and it dies on the day it acquires one.

There is a last question I find uncomfortable, which is whether a window should be logged at all. Logging creates a record, a record creates visibility, and visibility invites the assessment the whole thing exists to defer. That is the observation principle turning up inside the design of its own remedy, and I do not have a clean answer to it.

The Answer Guarantee

Chapter 9 established that futility does more damage than fear, lasts longer, and is harder to reverse, and that silence from a manager's side reads as futility from the other regardless of what was actually happening. So: every idea raised gets a response within a fixed number of days. Ten working days is the starting value. The response is the deliverable rather than the decision, and it includes no.

A no with a reason attached preserves the relationship and preserves the person's willingness to raise the next one, both of which are worth more than the idea usually was. Nothing at all teaches a lesson that lasts a career. Maria at twenty-six would have had her answer in ten days, and the answer would probably have been that the queue was not going to be resorted, and she would still be raising things at thirty-seven.

Reliability matters more than speed here, which is why I would rather see fifteen days kept than five days promised. An unkept guarantee is worse than no guarantee, because it demonstrates the thing it was built to disprove, to exactly the people who were willing to test it.

Two questions I cannot answer. Who owns the response when an idea travels upward past the first manager, which is where

most of them are lost. And how to track it without building a system that becomes its own small bureaucracy.

The Reversible Router

Chapter 5's one-way and two-way doors, turned into two questions asked before any approval process engages.

Is this reversible? Is it under the threshold?

Both yes, and it skips the process entirely. The router's only job is to stop two-way doors receiving one-way-door treatment, and it is triage rather than evaluation — it does not ask whether the idea is good, which is the question that has no useful answer in month zero anyway.

Who sets the threshold is a real question and the level decides whether anybody believes it. And the router can defeat itself faster than either of the others: a form to fill in proving that something is reversible is a gate, and a slower one than the gate it replaced. The only version I can defend is a written standing statement rather than a check somebody performs.

. . . .

Together they are a route, which is the thing part five of the instrument below is about to ask you to write down. The router decides whether a thing needs to be inside the machinery at all. The window covers the period during which it cannot defend itself. The guarantee is what happens when somebody brings you something regardless of any of this, which they will.

None of the three is new. Every good manager I have worked for has done versions of all of them, quietly, for people they knew well. What makes them a route rather than a favour is that they are written down, they have numbers against them, and they are available to somebody the manager has never had a reason to look at.

Where the evidence stops

The diagnosis in this book is better evidenced than the cure, and the gap is wide enough that I would rather open with it than have it found.

The title is Kotter's and the argument is not. John Kotter named the dual operating system in *Harvard Business Review* in 2012 and expanded it the following year in *Accelerate*, and anybody who has read him will recognise the phrase on this chapter's first page. His second system is a network: a volunteer coalition drawn from across the hierarchy, with a guiding structure and a strategic change agenda, running alongside the management hierarchy and working on initiatives the hierarchy is too slow to handle. What I have described is not that. It has no coalition, no volunteers, no agenda and nobody in it, because it is not a structure at all — it is a set of wiring rules about who may see what and who may spend what, applied to the hierarchy that already exists. Kotter's answer to organisational slowness is to add a second organisation. Mine is to stop the first one destroying what its own people build. I am borrowing his name for the shape of the problem, which is two things running at once at different speeds without one consuming the other, and departing from him on the remedy. It should also be said that his framework is a practitioner model rather than a research finding, and I have not found a controlled test of it in either direction.

Everything in Parts II and III rests on documented mechanisms — capital lumpiness, attention as a zero-sum resource, promotion for visible events, programme decay, the voice and silence literature, and Bernstein's controlled result. This chapter rests on a smaller and more contested body of

work, and on an assembly of it that is mine. Read the two at the same confidence and you will be over-trusting the second one.

Ambidexterity is supported and it is not settled. The idea that an organisation can run exploration and exploitation at once, and be better off for it, has meta-analytic support: Junni and colleagues pooled sixty-six studies in 2013 and found a positive and significant relationship with firm performance. Against that, Mathias published a meta-analysis in 2014 covering a hundred and seventeen studies and more than twenty-one thousand firms, and found that ambidexterity yielded *weaker* effects than a focus on exploration or exploitation alone — the opposite of the received position, and from a larger sample.

It is a book chapter rather than a journal article, and the replications of its result cluster in small and early-stage firms rather than in large ones, which is not the setting this book is about. That bounds it. It does not dismiss it, and nobody has shown the positive result holds cleanly in a manufacturer of the size I have been describing either.

Psychological safety is the robust part, and its own author supplies the caveat. Frazier and colleagues pooled a hundred and thirty-six samples and more than twenty-two thousand individuals and found psychological safety related to information sharing, learning behaviour and performance, with the sharing and learning links strongest. Edmondson's own framing is the one to keep: it is the engine, not the fuel. Necessary, and on its own insufficient, which is Chapter 9's fear-and-futility split arriving from the literature rather than from the argument.

The four boundaries were designed for an assembly line. Bernstein set them out for zones on a production floor in a Chinese contract manufacturer, with a defined standard method and operators deviating from it. Everything in the section above is my translation of them onto knowledge work and AI tool use,

and no part of that translation has been tested. The decision-rights-versus-improvement-rights distinction is the piece I am most confident in and the piece with the least behind it, and those two facts sitting together should make anybody uncomfortable, including me.

The token argument has no study behind it at all. That central control produces worse spending discipline than local ownership is an argument from the learning curve and from ordinary experience of how budgets behave. Nobody has run it. The line about approval costing more than the thing approved is an illustration, and it is flagged as one where it appears. If somebody measures this properly and finds that centralised AI procurement is cheaper per unit of value delivered, that section is wrong and the rest of the chapter is unaffected.

And channel-don't-ban rests on survey work. The observation that organisations succeeding with AI tend to formalise what employees were already doing, rather than start from a platform decision, comes from industry survey research rather than from a controlled comparison. It is consistent with everything else here, which is a reason to find it plausible and not a reason to treat it as demonstrated.

And the three mechanisms have numbers in them that nothing supports. Six weeks. Five hundred euros. Ten working days. Each of those is research turned into something a person can grant on a Tuesday, and the research is silent about where any of the dials should sit, because nobody has run the experiment that would tell you. I have set them where they look plausible to somebody who has sat in the meetings, and that is the whole of the justification. If somebody runs a three-week window and it works as well as a six-week one, the parameter is wrong and the mechanism is not, and I would rather that were found out quickly by people running them than defended slowly by me.

Here is the study that would settle it. Take one company with more than one comparable site. In one site, change the wiring and nothing else: write down the decision-rights and improvement-rights distinction, firewall any monitoring output from anybody who writes a review, put a small owned budget in the hands of the people doing the work, and set a protected period rule. Change nothing at the other site.

Then measure, over two years, the one number that matters — not tool usage, which will move for reasons that have nothing to do with any of this, but the rate at which employee-built tools are voluntarily shown to somebody with authority, and the rate at which they survive being shown. If that rate does not separate between the two sites, the resolution in this chapter is wrong and the diagnosis in the first eleven stands on its own.

The problem with that study is the problem with every study I have asked for in this book, and by now it should be recognisable. Announcing it is an observation event. The site that knows it is the treatment arm has already been changed by knowing, and the honest version needs the two arrangements to be ordinary facts about how each site runs rather than an experiment anybody has been told about — which means a longer, duller, more expensive design than anyone is likely to fund.

. . .

Surface Check

So — should Maria raise her hand on Tuesday?

No, and this is the chapter where that answer stops being about her. Every condition described above is a condition somebody else has to create. She cannot decouple improvement rights from evaluation. She cannot set a protected period, aggregate a report, or firewall anybody from a review, and there is no version of her raising her hand that produces any of it. Nine chapters of Surface Checks have said not yet, and what the not-yet was waiting for turns out to be a set of arrangements that nobody in that building has ever heard described.

And the room she would be walking into tomorrow has been framed since Monday by a document about unmanaged usage. Whatever she says about a private account and eighteen euros a month arrives in that room as a confession, and it will be received as one, by people who are not being unfair.

The tool is nine months old and it is ready. The plant is not.

INSTRUMENT

The Dual-Operating-System Canvas

For the person with a budget and an org chart, because everything in it requires one. An afternoon, and most of the answers are already in your systems rather than in your judgement.

Part one. Sort the last ten.

Write down the last ten changes of any size that happened in your area. For each, mark which physics it belongs to: irreducibly large — capital, coordination, all-or-nothing — or local, reversible and compounding.

Then mark, for each one, which route it actually travelled.

If everything went through the same route regardless of which column it sits in, you have one system. That is the finding, and it is the only part of this instrument that does not need an afternoon.

Part two. The Submarine Lifecycle Map.

Take anything in your area that somebody built for themselves — and if you cannot name one, that is Chapter 11's finding rather than evidence of absence. For each, place it on a line with four points: built, working, used by a second person, supported by something other than its author.

Most things stop between the second and third points and stay there for years.

For each one, write the next transition and who owns it. Not who should own it. Who does, today, by name — and if the answer is the builder, the thing has no adulthood available to it and will leave with them.

Part three. The two rights.

For one department, write two lists. What may a person change without asking anybody, because it affects only their own work? What requires a decision, because it changes somebody else's?

Most organisations have never written this down, which means the answer defaults to *everything requires a decision*, which is what people assume in the absence of being told otherwise.

Then circulate the first list. That single act is the largest thing in this instrument and it costs nothing, and you will find out within a month whether your area's fleet was submerged because of fear or because nobody had ever said.

Part four. The Token Budget.

Three numbers and a sentence.

1. The default — what may anybody in your area use, today, without asking, at what cost. If the answer is nothing, that is your first number.
2. The allowance — the amount a person may spend in a month on their own judgement, exhaustible, theirs to run out of.
3. The escalation — the route to the expensive tier, who approves it, and how long it takes.

Then the sentence: *what does it cost this organisation to approve one month of that allowance?* Include the meeting. Compare the two numbers and stop the exercise if they are the wrong way round, because everything else in this instrument is downstream of that comparison.

Part five. The route, and whether it exists.

Chapter 10 asked the person holding something to work out what would have to be true before they could show it. This is the same question from the other chair, and it is the one to answer last.

What is the shortest route in your area, from a working thing that one person built to a supported thing the organisation owns, that does not pass through a business case?

Write it as a sequence of steps with names against them. If you cannot, there isn't one, and the fleet in your area is not a culture problem.

Then: has anybody travelled that route in the last twelve months? Name them, and find out how long it took, and ask whether they would do it again.

The last question is the one that will tell you the most, and it is the only one in this instrument you cannot answer from your own records.

And if there is no route, the section above describes three that fit inside an ordinary manager's authority. Take one, put your

own numbers on it, and the question becomes answerable within a quarter.

Next: Chapter 13 — whether the technology that caused all of this is also the thing that ends it, and what Tuesday actually does to the choice Chapter 1 opened with.

The line is still in the deck. Nobody has to do anything for it to be read out.

Does AI Break the Pattern?

The meeting was at nine and he was in the room at ten to, which is not a habit of his, and the reason for it that morning was that he wanted to see the screen before there were people sitting in front of it.

Eleven were on the invitation and nine came. It is the room on the first floor with the long table and the window onto the yard, and the blind on that window has been broken at one end since before he arrived, so on a bright morning half the table gets the sun and the people who know sit on the other half. He sat on the sunny side, because he had come in early and it had not occurred to him, and then spent the first part of the meeting unable to read his own screen.

Six items. His was fourth.

Quality Systems went second. He had seen the deck before — not this version, but a version, in October, and probably one before that. The programme is nineteen months old now. The woman presenting it is competent and was not the person who scoped it, which she has never said in a meeting and does not have to, because everybody in that room can do the arithmetic between what the roadmap slide showed in October and what it

shows now. Two milestones have moved a quarter to the right. One has been renamed.

The roadmap slide is still very good.

His own item took four minutes. He said there was a discovery report now, monthly, that the first one had found sixty-one people on the site using generative AI tools through company machines with nothing behind any of it, and that he was proposing two things – a licensed platform, and a communication to line management. There were no names in the deck and he read none out. He said, because it is true, that this was not about catching anybody. The line on the slide said *establish visibility of unmanaged AI usage*, and he read it in the flat voice you use for a slide you wrote yourself and have already stopped hearing.

Nobody objected. Somebody from IT asked which platform and he said that was the next question rather than this one. Somebody asked whether the communication should come from him or from the plant manager, which is the sort of question that means a thing has already been agreed. It was minuted as an action with his name against it and a date in October.

He had thought about the whiteboard twice since Monday – once on the walk to the car, and once at about half past six that morning standing in his own kitchen – and both times he had got as far as understanding that a line about establishing visibility produces, in the ordinary course of events, a cloth. Both times the next thought did not arrive. There is no sentence available to him that means *leave that board alone*, there is no meeting at which such a sentence would be received, and he is not a man who says things in rooms without knowing what they cost. So the line stayed where it was and he read it out, and his four minutes ended, and the small unfinished feeling he had about it went wherever those go.

Item five was the one the invitation had been about. *Please come prepared to share what your team is working on* — a sentence somebody in the group function writes into a calendar entry with no idea what it does to the people who read it. It ran the way those items run. Two people described things that were being procured. One described a pilot that had also been described at the last one.

So he asked the room the ordinary question, because it was his item in everything but name and because the round was thinning out. Was anybody actually doing anything with any of this. Not in a pilot. In their week.

. . .

Maria had been in the room since five past nine, on the shaded side, third seat from the door, with a printed agenda in front of her that she had not written anything on.

What she felt when the question came was not fear, and it is worth being exact about that, because fear is the word that gets installed afterwards by whoever tells the story and it is the wrong one. She has been afraid in that building. At twenty-six, waiting on a decision about a sort order that never came, she was afraid for about three weeks and then stopped, and what replaced it has lasted eleven years and does not feel like anything at all. This was flatter than fear and much faster. It was the sensation of arriving at a calculation you have run so many times that you can no longer catch yourself running it.

She had decided on Sunday evening. Then again on Monday, and again in the corridor at two minutes past nine, and each time the decision came out the same and each time it presented itself as fresh.

So she said she had a way of doing the deviation reports now. That it was quicker. That it pulls the week together and gives her

a summary on a Friday morning, and that it saves her a fair amount of time.

Every word of that is true and none of it is the thing.

She did not say that it clusters by failure mode across suppliers rather than by supplier, which is the whole of what it does and the entire reason it works. She did not say that it caught the seal defect five weeks early in the spring, which sits on the record as an ordinary catch by an experienced engineer and has never been attached to anything. She did not say *my own account*, or *eighteen euros*, or *nine months*. None of that was concealed in the sense of a decision taken not to say it. There was a version of the sentence that fitted the room and a version that did not, and the one that fitted arrived first — the way it had arrived first in the canteen on Friday, in front of a colleague with no authority over her whatsoever.

It took about forty seconds.

The response was warm and it was not empty. Two people said versions of that sounding useful. The man from logistics, sitting almost opposite her, the one who made the remark about pallet labelling four years ago, said he had wondered whether something like that was possible. And the man with *excellence* in his title said it sounded like exactly the kind of thing they wanted more of, and that the licensed platform, once it was in, would let people do that sort of thing properly.

He meant it as encouragement, and it is encouragement, and he is right about the platform. Read from her chair it is also a description of what would happen to her tool if she put it in front of the thing he has just been given an action to buy — that it acquires a scope and a sponsor and somebody's name, which is Chapter 8 arriving in a sentence offered kindly across a table by a man trying to be helpful.

She said that would be good.

The item moved on. Somebody asked about a supplier portal. The meeting finished eight minutes early, which is the only thing anybody in it remembered about it by Thursday.

Nothing was suppressed in that room. Nobody was punished, nobody was overruled, nobody was so much as discouraging, and the two people who between them held both halves of the answer sat four seats apart for fifty-two minutes and neither of them did a single thing wrong. He asked the question this book says he should ask. She answered it truthfully. The answer was smaller than the thing, and he heard it at the size it arrived in, because there is no way for a person to hear a thing at a size other than the one it arrives in.

A reader who has come through twelve chapters has been braced for a meeting that goes badly, and this one does not go badly, and that is deliberate. The version where somebody gets punished is the version everybody already believes and it is the rarer one. Twelve chapters of machinery, and what the machinery produces at the end of it is fifty-two minutes in which nothing happens and nobody is at fault.

The tool ran on Friday morning, the way it has every Friday for nine months.

. . .

The question, asked properly

Chapter 2 ended on a question and then spent ten chapters not answering it, which was not an oversight. It could not be answered until the machinery had been laid out, because the answer turns on which parts of the machinery a technology can actually reach.

Is AI just the sixth wave of a pattern that has survived all five before it, or is it, for the first time, the thing that changes the pattern itself?

This is the most conceptual chapter in the book and I would rather say so at the top than have it inferred from the thinness of the citations. What follows is argument. The evidence in it has all appeared earlier and is doing a second job here; nothing new is being introduced that would settle anything, because there is nothing available that would. The honest description is that I am reasoning forward from a pattern, and reasoning forward from a pattern is the thing this book has spent thirteen chapters warning about when other people do it. So I will give the other side its strongest form first, and I will not be giving it a straw version, because I think one of its three moves is dangerous to my position and I have not seen anybody make it.

. . .

The case that it breaks

Three arguments, in ascending order of how much trouble they cause me.

The gate's price collapsed. Chapters 4 and 5 established that an increment does not fail a business case so much as fail to be worth one. The approval cycle has a fixed cost — the form, the numbers, the slot in a calendar, the fortnight of preparing to be asked a question — and the increment is smaller than that cost by construction, which is why the scale threshold discards things for being small rather than for being wrong. Now watch that arithmetic when the artefacts cost an afternoon instead of a fortnight. The case and the risk note write themselves, and the ratio that made the whole exercise irrational inverts. On this reading the submarine was never a response to hostility so much as a response to paperwork, and the paperwork just got cheap.

Adulthood got cheap with it. Chapter 12 said an improvement needs different conditions to survive being born and to survive

being twenty, and that the adulthood conditions are documentation, a second user, a place it lives and somebody accountable for it. Of those, documentation has always been paid for in the currency people have least of, which is why the continuous-improvement literature has spent decades watching local gains evaporate.

Maria's tool leaves in a laptop bag next spring because nobody wrote it down. Writing things down is the single task this technology is best at. A submarine-to-surface transition that used to cost a project could now cost a Thursday afternoon, and if that is true then the whole of Chapter 12 is a solution to a problem that is about to solve itself.

And the nursery may no longer be needed. This is the serious one. The entire justification for the submarine, from Criscuolo onward, is delayed assessment — the embryo gets to become what it is going to be before anybody judges it, and there is no path from the first version to the approvable version that runs through a meeting. That argument has a hidden premise, which is that the distance between the two versions is measured in months.

Collapse the cost of building and you collapse the distance. If the sixth version of a thing can exist by the end of a weekend, then the window during which the idea is unapprovable shrinks toward nothing, and the person can arrive at the defensible version before anybody has had the opportunity to judge the indefensible one. There is nothing to protect, because the vulnerable phase is over by Tuesday.

Put together, that is a coherent and optimistic account in which the mechanisms in Parts II and III remain entirely real and stop mattering, because the technology routes around them from below. I do not think it is right. I do think a reader could put this book down at the end of Part III, read the same

evidence, and arrive there, and I would not be able to call them careless.

. . .

The case that it is the sixth

Every previous wave lowered the cost of building too. That is what a technology wave is. The mainframe put computation inside the firm, the personal computer put it on the desk of the person with the problem, client-server put the data where they could reach it, the web removed the need to own any of it, and mobile removed the need to be anywhere in particular. Each was described at the time as the thing that would let people just get on with it.

The Durability Ledger from Chapter 2 is the instrument for exactly this, and it returns the same answer every time you run it: the shape is older than the tool, and the tool did not move it. Schön has every element of Maria's Tuesday in a paper published in 1963, when the newest of the five was a machine in a room most of the company was not allowed into.

Two of the three currencies do not touch the cost of building at all. Capital lumpiness is a property of plants, coordination and irreversibility, and a cheaper way to write a proposal does nothing to it — you still cannot half-retool a line, and the machinery built to guard that decision will still be the machinery a small reversible thing is routed through, because that is the only machinery there is. Careers select the same way regardless of what tools cost, because Chapter 8's mechanism runs on visibility of events and a six-hour tool is not an event.

And the third currency gets worse rather than better. Attention is zero-sum, which is the property that makes Bouquet and Birkinshaw's finding bite, and cheap building increases the number of things competing for a fixed pool. Ten

people who can each produce a small working thing in a weekend have not made it easier for any one of those things to be seen. They have made it harder, and they have done it to each other, while the large programme in the room retains every gram of structural weight it had before. Maria's tool was never rejected. It was never seen, and forty more like it in the same building does not improve the odds of anybody turning to look.

Territory does not care how good the case is. Chapter 4's second killer was never about evidence. An increment gets defended against because it sits in somebody's marked water, and the reaction to infringement is a reaction to the marking. A flawless business case produced in an afternoon buys a faster and better-documented no.

Then the asymmetry, which is the one that decides it for me.

Every previous wave lowered the cost of building without lowering the cost of watching. The personal computer handed an engineer a spreadsheet and handed her manager nothing comparable; reading her afternoon still took a person sitting down to read it, and no organisation has ever had that person spare. That asymmetry is the reason the underground survived five technologies, and what protected it was never policy but the cost of looking.

This wave is the first one to lower both, in the same purchase, in about two years, as a side effect of buying something else. And of the two capabilities that arrived together, only one of them has a budget holder, a governance argument and a line on a slide behind it. Chapter 11 is the whole of that story and I will not run it again. What matters here is what it does to the fork: the technology that makes the submarine nearly free is the first one that also makes hunting it feasible, and those two effects are not independent, because the second is a reaction to the first.

Which leaves the optimistic case's third move, and I think it fails for a reason that has nothing to do with speed. The nursery

was never protecting the idea from the calendar. It was protecting it from the observer. Chapter 10 established that observation does not read a value off the thing, it fixes the thing at whatever is legible on the day of looking, and legibility has almost nothing to do with maturity. Maria's tool on Tuesday morning was nine months old, worked, and had a catch on the record. She still described a faster report, in forty seconds, to a room full of people who wanted to hear something good.

Speed of building does not touch that. You can arrive at the sixth version by Sunday night and still have exactly the same conversation on Tuesday, because what shrinks the thing has nothing to do with its age and everything to do with being looked at by somebody who could act on what they see.

And the tool that clusters by failure mode did not come from nine months of building. It came from nine months of running against real reports, which is a different resource and one that does not compress.

. . .

Where I land

Neither branch, and the reason is more useful than either.

AI is the first of the six waves to change the cost on both sides of the disclosure decision at the same time. Building got nearly free, which is the optimistic case and it is real. Watching got nearly free, which is Chapter 11, and it is also real. Every previous wave moved one lever and the pattern absorbed it. This one moves both, in opposite directions, and the outcome is not determined by the technology, because the technology supplied both of the things that would determine it.

So the honest answer to Chapter 2's question is that the fork is real and the branch is not chosen by the tools. It is chosen by wiring, and the wiring is a small number of unglamorous

arrangements that Chapter 12 spent itself on: whether the report comes at function level or with names in an appendix, whether anybody who sees individual output also writes reviews, whether a person may change their own work without asking, whether there is a route from a working thing to a supported thing that does not pass through a business case.

That is a smaller claim than either side of the fork and it is stranger than both. For sixty years the persistence of this pattern was a property of the system — nobody chose it, nobody could have unchosen it, and the technology of the day was not capable of moving it in either direction. That is no longer the case.

The pattern's survival over the next decade is going to be a decision, and it is going to be taken mostly by people who will not experience themselves as taking one. It will be taken in the default configuration of a product that arrives with individual-level reporting switched on because that was the easier thing to build, and in the ninety minutes somebody spends writing a covering note to go with it. It will be taken by a man with a good record and no bad intent, reading out a line he wrote in forty minutes because he could not think what to put in its place.

This is my argument and not a finding, and the next section says how far it runs past the evidence.

. . .

The choice Chapter 1 opened with

I said at the start that we would come back to her and that by the end we would know what she should do, and that at that point neither of us did. So here is the closing of it, and it is not the shape I expected when I wrote that line.

Chapter 1 put the choice as a binary. Surface it and probably lose it, or hide it and definitely lose it, later, quietly. Two branches, one object, one decision, taken on a Tuesday.

Every part of that turned out to be wrong except the arithmetic.

There was never one object. Chapter 10 took that apart: the thing that exists in the dark and the thing that exists in the light are different objects, and the act of moving between them is what makes the second out of the first. And there was never a moment of decision, because the decision was made at twenty-six, in front of an alphabetical queue that is still alphabetical, and what happened on Tuesday morning was recall rather than deliberation — which Chapter 9 said in as many words, and which I did not fully understand the consequences of until I had to write the meeting.

What she actually did was both branches at once, and neither of them properly. She surfaced, in the only form available to her, which was a true sentence about a faster report. She stayed submerged, because the thing that makes it worth having was still on her laptop when she sat down. The plant received neither the risk nor the benefit. And nobody in that room could have told you which of the two things had just happened, including her.

Chapter 1 also predicted the ending and I did not notice at the time that it had. *It will not be rejected, because rejection would require somebody to look at it first. It will simply never be seen, which is a quieter fate and one that leaves no record at all.* Twelve chapters later, that is exactly what happened, and it happened in a meeting convened specifically to look, with the right question asked out loud by the right person.

So: should she have raised her hand?

She was right, and being right is not the same as being well served. Every term in her arithmetic was bent by machinery

nobody built on purpose, and the three things that make a hidden improvement survivable in daylight — a rate she can defend, a second person who has used it, an event it has already survived — she has one of, and two of them require somebody else. Nobody in that building is building them. Call none of that a failure of nerve or a failure of goodwill; it is what one operating system and no route produce in a building where everybody involved is reasonable.

She transfers in the spring. The tool goes with her, and the deviation reports start piling up again in the summer, and the seal defects go back to being caught five weeks later, and nobody connects any of it to anything because nobody knew there was anything to connect. There will be no incident and nothing will appear in a report. The plant will simply be slightly worse at something it was previously good at, for reasons that are not written down anywhere.

That is the default ending. It is happening now, in most large organisations, at a scale nobody can measure, and the only thing standing between it and a different ending is a small number of people deciding to change the wiring before they buy the instrument.

. . .

Where the evidence stops

This chapter is argument standing on evidence presented earlier, and the argument runs further past its supports here than anywhere else in the book. The whole of it should be read at a lower confidence than Parts II and III, and lower than Chapter 12, which was already the provisional half.

The six waves are a framing, not a dataset. Nobody has measured rates of hidden innovation across technology generations, because nobody could — the phenomenon is

invisible by construction and the baseline does not exist for any decade, including this one. The Durability Ledger is a habit for noticing that a problem is older than its name. It is not an instrument, it produces no number, and I have used it here as a structuring device rather than as evidence.

The fork is not falsifiable in the short run. Both branches predict a great deal of what is currently visible. Rising official adoption is consistent with the pattern breaking and equally consistent with the fleet relocating and the compliance number improving as the estate goes dark, which is Chapter 11's whole warning. Anybody who tells you in 2026 that the data already settles this is reading their preferred branch into an ambiguous signal, and that includes me if I ever do it.

The both-costs-fell asymmetry is mine and it is unquantified. The claim that this is the first wave to lower the cost of building and the cost of watching in the same purchase is an argument about relative magnitudes that I cannot put numbers to. Somebody could reasonably say that the web lowered the cost of watching too, through logging and email retention, and that the difference here is one of degree. I think the degree is large enough to be a difference in kind, because summarising judgement work at scale was not previously possible at any price. But that is a judgement about a threshold and I have not measured where the threshold is.

And the optimistic case may be right. Its third move is the one I would attack this chapter with. If building becomes cheap enough that the version shown in month zero is already the good version, the observation principle loses most of its bite, because there is no better hidden version for the visible one to be a diminished copy of. My answer is that Maria's tool got good by being run against nine months of real reports rather than by being built for nine months, and that running time does not compress the way building time does. That answer is an

argument about where the value in an increment comes from. It is testable and it has not been tested, and if it is wrong then the pessimistic half of this chapter goes with it.

Here is the study that would settle it, and it is the one I would most like to see run. Take a cohort of employee-built AI tools now, identified in a way that does not require anybody to confess to anything — through a voluntary registry with no evaluative consequence, or through a third party with genuine anonymity. Record for each one the date it was built.

Then, five years from now, measure two proportions: how many were voluntarily shown to somebody with authority within their first year, and how many survived being shown in a recognisable form. If both rise substantially while the cost of building stays low, the pattern broke, the optimistic case was right, and this book is a period piece about the last years of a sixty-year phenomenon. If they hold where they are, it is the sixth wave.

The problem with that study is the one that has attached itself to every study I have asked for in this book, and by now it should be familiar. There is no baseline for the first number, in this decade or any of the previous five, and there is no baseline because the thing being measured hides from measurement, and the reason it hides from measurement is the subject of the book. A phenomenon that cannot be counted is real anyway; it is only unfalsifiable at the speed anybody would like.

. . . .

Surface Check

So — should Maria raise her hand on Tuesday?

She did. Not in the form the question has been imagining for twelve chapters — no revelation and no risk — but in the only form her plant made available to her, which was a true sentence forty seconds long about a faster report. Nothing was punished. Nothing was adopted. The question has been answered by being asked in a room, and the answer is that it was never a question about her.

Every Surface Check since Chapter 1 has said not yet, and every one of them was right, and what they were waiting for was three things, two of which require somebody else. Nobody in that plant is building them, and nobody will before the spring.

This is the last one. The device has done its work, and the honest thing to say at the end of it is that the person it kept returning to was never the one who could change the answer.

INSTRUMENT

The Monday Move

The last instrument, and unlike every other one in this book it does not ask you to audit anything. One action each, for the three people who read this, and each of them small enough to do without asking anybody.

That constraint is not modesty. A move that requires a programme has already failed the argument it comes from, and there is no version of this book that ends by asking you to get something approved.

If you have built something.

Do not surface it. Get a second user.

One person, by name, this week — somebody who will run the thing themselves rather than watch you run it. Of the three things that make a hidden improvement survivable in daylight, it is the only one available to you without anybody's permission, and it is the one that converts a tool that leaves in a laptop bag into a tool with a witness.

Be clear about the price, because there is one and I am not going to pretend otherwise. A witness is also a person who can mention it in a room you are not in, which is precisely what happened to Maria on a Friday in the canteen. What you are buying with that risk is the only thing that makes disclosure survivable later, and the sequence matters: second user first, disclosure afterwards, and not the other way round.

Then a date. Not a verdict — Chapter 10 already made that point and this is the practical end of it. Write down when you will ask yourself the question again, and what would have to have changed by then, and put it somewhere that is not a company device.

If people report to you.

Write the two lists. One page, one department, an hour of your time.

What may somebody in your team change without asking anybody, because it affects only their own work? What requires a decision, because it changes somebody else's? Most organisations have never written this down, which means the answer defaults to everything requiring a decision, because that is what people assume when nobody has told them otherwise.

Then circulate the first list. That single act costs nothing, needs no approval, and within a month you will know whether the fleet in your area was submerged because of fear or because nobody had ever said.

If you only do one thing and it is not that, do Chapter 9's instead: the last five ideas anybody brought you, what happened to each,

and — the part that matters — whether the person found out. Silence from your side reads as futility from theirs, whatever was actually going on, and futility is the one that does not reverse.

If you sign things.

Before the instrument goes live, one email.

Individual-level output from this data does not go to anybody who writes, contributes to or signs off a performance review. The report arrives at function or site level. Names are not in the appendix.

That is three sentences in a configuration document, it costs nothing, it survives being made free, and it is the entire difference between the developmental case and the evaluative one — which is to say it is Chapter 11 reduced to something you can send before lunch. The governance question you actually have to answer survives aggregation completely. The appendix is the part that does not, and the appendix is the part that gets read.

Then the move that creates something rather than preventing something, because an email that stops a harm still leaves your area with no route. Grant one Protected Window this month. Six weeks, a cap you can sign off without asking anybody, exempt from assessment for the duration, and one conversation at the end — Chapter 12 has the parameters and the reasons they are the weakest thing in this book.

The conversation at the end is where the value is and it is the part to get right. Ask what they learned. Then wait, and do not fill the silence.

And all three of you.

Leapfrog closes on placing yourself on the roadmap, and it is the right instruction to end this one on as well, with a different roadmap in mind. You are one of those three people. The move for your seat takes under an hour, requires nobody's permission, and is not the move for either of the other two — which is most

of why this problem has survived sixty years of everybody agreeing that somebody ought to do something about it.

Pick yours. Put a date on it.

The tool ran again on Friday morning, the way it has every Friday for nine months. Maria transfers in the spring.

CONCLUSION

What To Do About It

A book that has spent thirteen chapters warning against confident inference from incomplete evidence owes the reader a straight account of its own footing before it starts issuing instructions.

So here is the account, and then here is the instruction, and the second does not depend on the first being settled.

. . .

What this is

A theory. Not a law and not a finding — an explanation that fits a large body of published research, a much smaller body of direct measurement, and one pattern that has held for sixty years, and which makes predictions that could turn out to be wrong.

Four tiers. They are set out in one place here because thirteen separate sets of caveats are easy to lose track of, and the total is more useful than the instalments.

Established, and not by me. Bootleg effort raises innovative performance, and the mechanism is delayed assessment rather than avoided assessment. Monitoring produces no reliable performance gain, across ninety-four studies and more than

twenty-three thousand people. Expecting to be evaluated degrades creative work, which Amabile established in a laboratory in 1990. Observation degrades performance on a production line, which is Bernstein's controlled result with a curtain in it. Attention at headquarters is allocated by structural weight and there is only so much of it. Voice extinguishes through futility faster and far more durably than through fear. And capital processes are built to guard irreversible decisions, then route reversible ones through the same machinery because it is the only machinery there is. None of that is in dispute and none of it is mine.

Mine, and well supported. That those mechanisms converge — that capital, attention and career incentives all bend the same way against the same class of small improvement. No single literature says this, because no single literature is looking at all three at once. And the observation principle: that one effect operates at the scale of the idea, of the programme and of the person. That is a synthesis across three bodies of work rather than a finding inside any of them, and I have been careful to say so each time it appears.

Inferred from published models, and not measured. That concealment relocates rather than stops when the cost of visibility rises. An enforcement record from another industry shows that behaviour under prohibition; nobody has measured it under monitoring, and those are different stimuli. Also the two-conditions frame — hide to survive birth, surface to survive adulthood — which is a reading of what the birth-condition research implies rather than something anybody has tested.

Conjecture, with a test attached. That AI is the first wave to move the cost of building and the cost of watching in the same purchase, and that the branch we end up on is therefore decided by wiring. Chapter 13 sets out the study that would settle it and names the result that would falsify it. Nobody has run it.

That is the whole of the footing. Somewhere between a strong hypothesis and a working theory, and closer to the first than I would like it to be.

. . . .

Why you should act on it anyway

The obvious objection is that a theory in that condition does not license instructions.

It licenses some and not others, and the line falls exactly where this book has been putting it since Chapter 5. An action that is cheap and reversible does not need proof behind it. It needs to be plausibly useful, cheaply abandoned, and harmless if the theory turns out to be wrong. That is a two-way door, and the correct treatment of a two-way door is to walk through it and find out, because the finding out is cheaper than the deliberating. An action that is expensive and hard to undo needs a great deal more, because you only get to be wrong about it once.

Everything I am about to ask you to do is a two-way door. None of it costs real money, none of it needs approval above a single manager, and all of it can be stopped inside a week with nothing lost but the week.

The thing I am asking you not to do is the one-way door. Buying an instrument and pointing it at your own people is expensive, visible and effectively permanent in its effects — the people who went quiet do not come back when the policy is withdrawn, which is Chapter 9's mechanism and among the better-established things in this book.

I want to name that asymmetry before somebody else does. I am asking you to take small reversible actions on incomplete evidence while asking you not to take one large irreversible action on the same incomplete evidence. That is not

inconsistency. It is what a two-way door is for, and if everything else in this book turns out to be wrong, that second half is still the right way to treat a decision of that shape.

. . . .

First, the thing not to do

I have just told you that something valuable is hidden inside your organisation and roughly where it lives. The obvious use of that information is the worst move available to you.

Do not go and find it.

Not with a discovery report, not with a licensed platform that logs individual prompts, not with a well-meant survey asking people what they have built. What you find will be the version that survived being found, you will lose the rest of it permanently, and the compliance number on your report will improve the entire time this is happening. That is Chapter 11 and it is the sentence I would keep if I could only keep one.

. . . .

Then, in this order

Five things. The order is part of the argument, and each step is cheaper than the one that follows it.

One. Take the instrument off the table. If a discovery or monitoring capability is on a roadmap in your area, the first move is to stop it, and the second is to find out what question it was bought to answer — because there is usually a real one underneath, and it usually has an aggregate answer. How much of this is happening, in which functions, on what data. Every one of those survives being answered at function level with no names anywhere near it.

Two. Wire it before you buy anything. If something is coming regardless of what you want, then it arrives configured, and the configuration is the whole game. Individual-level output goes to nobody who writes, contributes to or signs off a performance review. Reporting lands at function or site level. No names in the appendix. Three sentences in a configuration document, sent before lunch, and it is the entire difference between a developmental instrument and an evaluative one.

Three. Open one route. This is the same act at three altitudes, and you do the version your seat allows. If you built something, get one other person using it — the only survivability condition available without anybody's permission. If people report to you, write down what somebody may change without asking and circulate it, which most organisations have never done, so the answer defaults to nothing. If you sign things, grant one Protected Window: six weeks, a cap you can authorise alone, exempt from assessment, one conversation at the end.

Four. Answer people. Ten working days, every time, including no. The response is the deliverable and the decision is secondary, because a no with a reason attached preserves somebody's willingness to raise the next thing and silence does not. This is the cheapest item on the list and the one with the longest reach, and it is the one most likely to be agreed to and then not done.

Five. Then look — and only by invitation. After the other four, and not before. Ask one person who is not on any list and does not report to you to show you something they made for themselves. Then wait, and do not fill the silence. What comes back will be smaller than the thing, because Chapter 10 says it must be, and it will still be more than any instrument would have found you.

. . . .

If I am wrong

Suppose the optimistic branch is the right one and the pattern is dissolving on its own, because building got cheap enough that nobody needs to hide anything any more.

Count what following this list will have cost you. Some people got answers within ten days. A few six-week windows ran, and some of them produced something and some produced nothing, which is what a window is for. Somebody wrote down what their team may change without asking. And an instrument that would have been unnecessary under that branch anyway did not get bought.

There is no version of that branch in which any of it was a mistake, which is the property I was selecting for when I chose these five and not others.

And if the pattern holds, you will have built the only thing that makes any difference to it, which is a route — one path from a working thing somebody made on a Thursday to a thing the organisation owns, that does not run through a business case. Most large organisations do not have one. Nobody has ever been asked to build one, because nobody has been able to see what it was for.

You can see what it is for now. That is the whole of what this book was written to do.

AFTERWORD

A Note on How This Was Made

The book is finished and Maria is still sitting in it, so let me come out from behind her for a moment.

The Conclusion has already told you what to do, so this is the other thing — where the rest of this lives, and one hope I can only state rather than ask for.

The hope first. I would like you to be unable to un-see this now. The next roadmap slide you are shown, and the next person on your team who is faster than their process should allow — I would like both of those to look different to you than they did a fortnight ago. If that is the only thing this book does, it was worth writing.

. . . .

Where the rest of it is

submerged.lerias.org has the research. All of it, including the studies that cut against me and the red team I ran against my own argument. Almost nobody publishes their sources and I have never understood why. If you want to check whether I have handled a finding fairly, you should be able to do that without my help.

The instruments from each chapter are there in a form you can use, rather than in a form you have to retype out of a book.

The mechanisms have their own page, and that matters more than it sounds. Chapter 12 puts numbers in your hands — six weeks, five hundred euros, ten working days — and says plainly that they are starting values with nothing behind them. A book cannot be corrected after it is published. A practice can. So the parameters live on the site with a version number against them, and they will move as people run them and tell me what happened.

Which is the one thing I would like in return. If you grant a window, tell me how it went, including if it went nowhere — especially then. I am not asking for that as a marketing device; it is the only mechanism by which any of this gets less wrong.

leapfrog.lerias.org is the other half, and it is free too. This book is the why and the when. That one is the how. If you have finished this and the problem you now have is that you need to actually ship something, start there.

This book is free as well, under a licence that lets you pass it on as long as you do not sell it. There is no list to join and there is nothing for sale, which is partly a principle and partly the terms under which I am allowed to write it at all. If you think somebody in your building should read it, send it to them. That is the whole of the distribution strategy.

. . .

One last thing

A note on how this was made, since a book about AI-assisted work ought to be clear about its own.

I wrote it on evenings and weekends over several months, on my own time and my own account. I used AI throughout — to assemble and cross-check research, to validate claims against

their sources, to draft and redraft passages, to review chapters for the things I had stopped being able to see, and to argue ideas into shape before I committed to any of them.

The argument is mine. So is every decision about what this book claims, what it refuses to claim, and where the evidence stops. Nothing went in because a tool produced it; a good deal came out because I did not believe it.

What the technology changed was the scale of what one person could attempt in the hours available. That happens to be the subject of this book, and I have no interest in pretending I am standing outside it.

That's it. Have fun.

APPENDIX

The Promises Ledger

Every chapter of this book makes claims it does not finish, and points forward to the chapter that finishes them. While it was being written, each chapter carried a running ledger of those debts at its head. That scaffolding has been taken out of the chapters, where it sat above the opening scene and told you what you were about to feel. It is collected here instead, so that a reader who wants to check whether the book paid what it promised can do it in one place, and in one pass.

Read the middle column as the debt and the right-hand column as the settlement.

PROMISE	OPENED	PAID
The pattern is permanent, and decades old	Ch 1	Ch 2
Hidden work is better, because assessment is delayed	Ch 1	Ch 1; mechanism, Ch 6
The hollowing-out set the stage	Ch 2	Ch 3
Every mechanism is individually reasonable	Ch 2	Part II, closing Ch 6
Maria's rival programme: eighteen months, two consultancies	Ch 1	Ch 5
The buried increment is worth something	Ch 5	Ch 6
Companies tried to domesticate the underground, and failed	Ch 2	Ch 7
Protection is itself a programme, and programmes decay	Ch 7	Ch 7, again Ch 9
Why nobody has fixed it	Ch 1	Ch 7–9
Escalation of commitment arrives with its warning attached	Ch 2	Ch 8
Careers are the third currency	Ch 5	Ch 8
The three-currency convergence	Ch 2	Ch 5, closing Ch 8

PROMISE	OPENED	PAID
Surfacing can destroy the thing surfaced – the curtain	Ch 1	Ch 10
The observation principle	Ch 6	Named Ch 10, spent Ch 11
The hunt backfires	Foreword	Ch 11
The legitimate reasons to look	Ch 10	Ch 12
Two conditions, not one decision	Ch 10	Ch 12
What an executive actually grants	Ch 11	Ch 12
Does AI break the pattern	Ch 2	Ch 13
The choice Chapter 1 opened with	Ch 1	Closed, Ch 13

Three faults were found in the per-chapter ledgers while they were being consolidated, and are corrected above. The three-currency convergence was marked as closing in Chapter 8 and then as still accruing in Chapter 9. *Protection is itself a programme* was opened in Chapter 7 and never closed anywhere. And *the domestication attempts fail*, opened under that name in Chapter 3, came back in Chapter 7 under another. Twelve copies of a nearly identical list is twelve places for a debt to go missing, which is a reasonable argument for keeping one.

ABOUT

The Author, and the Companion Volume

Hugo Lerias has spent his career inside the companies that built the modern internet — the edge era at Akamai, the early cloud at Amazon Web Services, and now AI-enabled engineering in manufacturing. He has never worked on a plant floor, and nothing in this book rests on his having done so. He writes about what large organisations do to small good ideas, which is a subject he did not choose so much as accumulate.

He is also the author of ***Leapfrog: A Corporate Engineer's Field Guide to Delivering AI on the Cloud*** — free at leapfrog.lerias.org. The two books share a spine: AI has revoked the corporate excuse for deferral. *Leapfrog* is the engineering answer to *how*. *Submerged* is the organisational answer to *why* and *when*.

Both are free, both are licensed CC BY-NC-ND 4.0, and neither has ever been for sale.

submerged.lerias.org