

BEYOND THE PILOT

The Last Mile of AI Adoption: Why Enterprise Resistance Is Rational

On the gap between what the tools can do and what organisations will let them do — and why the friction in between is a signal, not a defect.

Giovanni Leonardi

April 2026



Contents

- 1. Executive Summary** 3
- 2. The Pilot That Never Became a Practice** 3
- 3. The Comfortable Story** 4
- 4. Resistance as Signal** 5
 - 4.1 The accountability asymmetry 5
 - 4.2 The verification tax 6
 - 4.3 The tacit-knowledge frontier 6
 - 4.4 The measurement inversion 7
 - 4.5 The buyer and the user 7
- 5. The Strongest Objection** 7
- 6. Making Adoption the Rational Choice** 8

The Last Mile of AI Adoption: Why Enterprise Resistance Is Rational

Transformation · No. 369
April 2026

“The uncomfortable truth of the last mile is that the resistance has usually done the analysis the programme skipped.”

1. Executive Summary

For two years the enterprise conversation about artificial intelligence has circled a single, stubborn observation: the pilots succeed and the rollouts stall. A capability that dazzles in a demonstration and clears its business case in a controlled trial somehow fails to cross into the daily work of the organisation. The received explanation is that people resist change — that the last mile of adoption is a psychological problem, to be dissolved with communication, training, and enough executive sponsorship.

This essay argues something close to the opposite. The resistance that strands enterprise AI in its last mile is, for the most part, rational. It is not fear wearing the costume of prudence; it is an accurate reading of how accountability, risk, and reward are actually distributed when a probabilistic tool is dropped into a process that was built for accountable people. Once you look at the incentives the worker is actually facing — rather than the ones the transformation deck assumes — declining the tool, hedging with it, or quietly working around it is frequently the correct individual choice, even where the aggregate case for adoption is genuine.

The argument runs in three moves. First, that the comfortable story — resistance as a deficit to be overcome — is seductive precisely because it flatters everyone who is already convinced. Second, that beneath it sit five structural forces that make resistance a rational response rather than an emotional one: an accountability asymmetry, a verification tax, a tacit-knowledge frontier, a measurement inversion, and a split between the people who buy the tool and the people who must live with it. Third, that the practical consequence is not to argue people out of their rationality but to change the structure so that adoption becomes the rational choice. Read this way, the friction of the last mile stops being an obstacle to be steamrolled and becomes the most useful signal an organisation has about where its AI actually pays off — and where it merely appears to.

2. The Pilot That Never Became a Practice

There is a moment, familiar now to anyone who has sat through a wave of these programmes, when a proof of concept goes gently cold. The pilot was a success — everyone agrees it was a success. The tool did in seconds what used to take an afternoon. The slides showed the time saved, the satisfaction scores, the enthusiastic quotes from the volunteers who ran it. A licence was purchased for the whole department. And then, six months on, the usage dashboard tells an awkward story: a spike in the first fortnight, a long decline, and a stubborn core of people who logged in twice and never came back.

The programme is not failing loudly. Nobody has refused. There is no petition, no confrontation, no article in the trade press about workers downing tools. What there is instead is a quiet, distributed, almost invisible non-adoption — the organisational equivalent of a polite nod followed by carrying on exactly as before. In the reviews, this gets a name. It is called change resistance, and the name does a great deal of work, because it quietly settles the question of who is being unreasonable.

I have watched this pattern repeat across sectors that have nothing else in common, and the striking thing is how consistently the diagnosis arrives before the investigation. The tool is assumed to be good; the business case is assumed to be sound; therefore the only variable left to explain the shortfall is the people. And so the second wave of the programme is built almost entirely around them — more training, more champions, more mandatory-by-a-different-name targets, a communications campaign about embracing the future. The tooling is treated as settled. The humans are treated as the problem.

What almost never happens, in that second wave, is the question that a curious outsider would ask first: what if the people who stopped using it are right?

3. The Comfortable Story

It is worth being fair to the received view, because it is not stupid and it is not always wrong. People do resist change. Status-quo bias is real; loss aversion is real; there is genuine anxiety, in a period when the tools are explicitly sold on their ability to do what knowledge workers are paid to do, about what widespread adoption means for one's own standing. Some resistance is exactly what the change-management literature says it is: discomfort, inertia, and the ordinary human preference for the devil one knows.

The trouble is not that this story is false. The trouble is that it is comfortable — and comfortable in a specific, self-serving direction. It locates the problem entirely in the population that has the least power in the transaction and the least authorship of the decision. It requires no revisiting of the business case, no admission that the pilot measured the wrong thing, no uncomfortable conversation about whether the tool is actually ready for the messiest sixty per cent of the work. It lets the people who chose the tool keep believing they chose well, and reframes every piece of contrary evidence as a maturity problem in someone else.

The great convenience of "change resistance" as a diagnosis is that it can never be disconfirmed by the behaviour it describes. If people adopt, the tool was good. If people resist, they were afraid. The tool is protected either way – which is precisely why the diagnosis should be distrusted.

This is where the essay wants to plant its flag. Treating all resistance as irrational is not a neutral analytical error; it is an expensive one, because it throws away information. When a large number of experienced practitioners independently decline to use a tool for a large fraction of their real work, that is not noise to be managed. It is a distributed experiment, run for free, by the people who know the work best, on the question the organisation most needs answered: where does this actually help, and where does it merely appear to? The last mile is not the place where good decisions meet human weakness. It is the place where the pilot's optimistic assumptions finally collide with the real texture of the work – and lose.

4. Resistance as Signal

The reframe I want to propose is simple to state and awkward to act on. The resistance is data. Not a universal veto, not a proof that the sceptics are right about everything, but a signal with real information in it – a running, task-by-task verdict from the people closest to the work about where the tool's value survives contact with reality.

To take that seriously, you have to explain the mechanism. It is not enough to assert that resistance is rational; the whole burden of the argument is to show why a sensible person, acting in good faith and wanting to do good work, would decline a tool that demonstrably saves time. Below are the five forces I keep returning to. None of them is psychological. All of them would still operate on a perfectly rational, entirely unafraid worker – which is the point.

4.1 The accountability asymmetry

Consider what actually happens when an assistant drafts a piece of work and a person puts their name to it. If the work is right, the productivity gain is diffuse: it shows up in the aggregate, in the dashboard, in the executive's narrative about transformation. If the work is wrong – a figure subtly off, a clause that does not mean what it appears to mean, a recommendation that is confidently and plausibly incorrect – the accountability is not diffuse at all. It lands, entire, on the individual who accepted the output.

This is the deepest of the structural forces, and the least discussed. The tool socialises the upside and privatises the downside. The organisation captures the speed; the individual retains the liability. Asked to "just use the tool to go faster," a rational worker correctly perceives that they are being invited to take personal responsibility for the errors of a system they did not build, cannot fully inspect, and are not permitted to switch off when it matters most. Declining is not timidity. It is an accurate pricing of asymmetric risk. Any incentive scheme that leaves the error on the human while handing the speed to the enterprise will produce exactly the caution it then labels as resistance.

4.2 The verification tax

The promise of the tool is that it does the work so you need not. The reality, across a great many knowledge tasks, is subtler and less flattering: the tool produces a draft quickly, but bringing that draft to the standard the work actually requires — checking every figure, every reference, the one sentence in twenty that is quietly wrong — carries its own substantial cost. Call it the verification tax.

Where verification is cheap relative to production, the tool is transformative and adoption is effortless; nobody resists a spell-checker. Where verification approaches the cost of production, the calculus inverts. Consider a common pattern: a drafting task that once took an experienced analyst forty minutes. The assistant returns a serviceable draft in two. But the draft must be brought to a standard a client will see, and verifying unfamiliar work is often harder than producing familiar work, because the reviewer must reconstruct reasoning they did not perform. Suppose that check takes thirty-five minutes. The net saving is real — three minutes — but small, and it now arrives bundled with a new and worse failure mode: the plausible error that glides past a tired reviewer precisely because it looks like competent work.

Task characteristic	Verification cheap	Verification dear
Error is obvious when present	Yes	No
Reviewer must reconstruct reasoning	No	Yes
Consequence of a missed error	Low	High
Rational response to the tool	Adopt eagerly	Adopt selectively or decline

A worker who embraces the tool for the first kind of task and refuses it for the second is not exhibiting inconsistent "adoption maturity." They are performing an intelligent, task-by-task cost calculation that the organisation's blunt adoption metric cannot see. Counting logins, the enterprise reads a shrewd allocation of the tool to where it pays as a failure to adopt.

4.3 The tacit-knowledge frontier

Every pilot is run on the legible slice of the work — the clean cases, the well-documented process, the forty per cent that can be written down. The last mile is the other sixty: the exceptions, the undocumented rules, the "we do not do it that way here, for a reason nobody ever wrote down." The knowledge that governs the hard cases does not live in the systems the model can see. It lives in people, accreted over years, and much of it has never been made explicit because in a human-to-human organisation it never needed to be.

When an experienced practitioner resists applying the tool to the difficult residue of their work, they are very often exercising precisely the tacit judgement the tool lacks — the knowledge that this customer is a special case, that this number always looks wrong and always turns out to be right, that this apparently routine request is the visible edge of a problem the model has no way of knowing about. The resistance is the tacit-knowledge holder's correct assessment that the tool does not know what they know. Override it

wholesale and you do not cross the last mile; you simply move the errors somewhere less visible.

4.4 The measurement inversion

Organisations instrument what is easy to instrument. Throughput, cycle time, volume, tickets closed — the legible surface of the work. AI is very good at the legible surface, and so it accelerates exactly the part that was already measured, while the illegible part — judgement, exception-handling, the relationship that stops a small problem becoming a large one — becomes a larger share of what actually matters and remains stubbornly unmeasured.

Now watch the perverse result. The worker who slows down to catch the subtle error the assistant introduced is penalised by the very dashboard that is simultaneously celebrating the assistant's speed. The metric rewards the throughput the tool provides and is blind to the judgement the human is quietly supplying to keep the throughput safe. Faced with that, the rational response is not to argue with the dashboard. It is to give the dashboard what it wants and protect the real work in the margins — which, from above, looks like enthusiastic adoption sitting suspiciously alongside a slow drift in quality that nobody can quite attribute.

4.5 The buyer and the user

The final force is the one that ties the others together. The tool is chosen by people whose incentive is the transformation narrative, and used by people whose incentive is not to be the one holding the error when it surfaces. These two incentives are not aligned, and the last mile is precisely where they collide.

The buyer is rewarded for the story of adoption — the rollout, the licence count, the case study. The user is rewarded, or punished, by the consequences of specific outputs in specific situations. When the buyer mandates the tool and the user hedges against it, this is not a hierarchy failing to communicate. It is two rational actors optimising for two different, genuinely misaligned objectives. No amount of change communication closes a gap that is structural rather than informational, because the user already understands the message perfectly well. They simply have different — and equally rational — reasons to act.

5. The Strongest Objection

An honest essay has to meet the best version of the counter-argument, not a convenient weakling. Here it is, at full strength.

Much of what I have dignified as "rational resistance," a critic will say, is simply rationalisation — the sophisticated professional's gift for producing principled-sounding reasons for the thing they were going to do anyway, which is protect their own comfort, their own status, and their own scarce expertise from a tool that threatens all three. Every one of my five forces, the objection continues, can be recruited as an excuse. And there is harder evidence still: in this very period, the organisations that hung back waiting for the tools to be perfect, waiting for the accountability to be clean and the verification

to be free, have in several cases been overtaken by less cautious competitors who tolerated the imperfections, absorbed some errors, and moved. If caution were always rational, prudence would always win. It has not. Sometimes the resisters were simply wrong, and the cost of their rational-sounding hesitation was the market.

This objection is strong, and I do not think it can be waved away. Two things must be conceded to it plainly. Resistance is not always rational; some of it is exactly the fearful turf-protection the change literature describes, and the language of structural risk is indeed available as a cloak for it. And caution is not free; there is a real and sometimes fatal cost to waiting, and the organisations that treated every one of these forces as a reason to do nothing frequently deserved what happened to them.

But concede both and the argument still stands, because it was never that all resistance is rational — only that treating all of it as irrational is the more expensive of the two available mistakes. The firms that were overtaken were not, on the whole, punished for reading their workers' resistance too carefully. They were punished for the opposite failure — for organisational paralysis, for using caution as a synonym for inaction, for hearing "this tool is not ready for the hard cases" and concluding "therefore we will change nothing." The lesson of their defeat is not ignore the resistance. It is act on what the resistance is telling you: deploy hard where verification is cheap and accountability is clean, hold back where it is not, and spend the effort you would have wasted on persuasion campaigns fixing the structure instead. The winners were not the ones who overrode their sceptics. They were the ones who read them as instruments — and moved.

6. Making Adoption the Rational Choice

If the diagnosis is structural, then so is the remedy. You do not cross the last mile by persuading people to be less rational. You cross it by changing the structure so that using the tool well becomes the individually correct thing to do. Four consequences follow directly from the five forces, and none of them is a communications exercise.

1. Reallocate the accountability. If the enterprise wants the speed, the enterprise must hold more of the risk. That means explicit, institutional ownership of tool-assisted error — review structures, sign-off that is shared rather than dumped on the last person to touch the work, and an honest acceptance that a probabilistic tool will sometimes be wrong and the organisation, not the individual, will absorb it. The moment the person is no longer holding uninsurable risk alone, a great deal of "resistance" evaporates, because it was never irrational in the first place.
2. Measure the verification tax, not the logins. Adoption metrics that count usage are measuring the wrong thing. The number that matters is the net cost of the whole task including verification, and it must be measured task by task, honestly, including the tasks where the answer is that the tool does not yet pay. An organisation that knows its real verification tax can deploy the tool precisely where it wins and stop pretending about where it does not.
3. Put the tacit-knowledge holders in the design seat. The people resisting hardest are frequently the ones who know most about the difficult residue. They are not the obstacle to the last mile; they are the map of it. Bringing them into the design of where and how the tool is used — as authorities rather than as adoption targets —

converts the most valuable resistance into the most valuable design input.

4. Change the metrics before you change the tools. As long as the dashboard rewards throughput and is blind to judgement, it will punish exactly the behaviour that makes AI safe to use, and reward exactly the behaviour that makes it dangerous. Fixing the measurement inversion is unglamorous and slow, and it is the single change most likely to make honest adoption rational for the individual.

“The last mile is not crossed by arguing people out of their rationality. It is crossed by making adoption the rational choice — and until it is, the friction is not the enemy of the strategy. It is the most honest feedback the strategy will ever get.”

None of this is an argument against the tools. The capability is real, the gains where they land are real, and the organisations that never move will be overtaken by the ones that do. It is an argument against a particular and costly misreading — the reflex that hears hesitation and concludes weakness, that treats the people closest to the work as the friction in the plan rather than the sensor in the system. The uncomfortable truth of the last mile is that the resistance has usually done the analysis the programme skipped. The organisations that will cross it first are not the ones that shout loudest about embracing the future. They are the ones that stop long enough to ask why their most experienced people are holding back — and are willing to hear that the answer is not fear, but arithmetic.