

BEYOND THE WORKING MODEL

The Part Everyone Skips: Why AI Adoption Stalls Long After the Model Works

On the distance between a model that works and an organisation that has changed — and why we keep mistaking the first for the second

Giovanni Leonardi

September 2022



Contents

1. Executive Summary	3
2. A metre from the finish line	3
3. The economics of the skipped step	4
4. What "the human side" actually is	5
5. The objection worth taking seriously	6
6. Speed, trust, and the tension no method dissolves	6
7. What taking it seriously looks like	7
8. The mirror	8

The Part Everyone Skips: Why AI Adoption Stalls Long After the Model Works

Transformation · No. 286

Originally written September 2022 | Re-edited and re-rendered for the WhereBeyond Edition,
2025

*"The work everyone skips is not a phase of the project. It
is the project."*

1. Executive Summary

Most organisations that adopt artificial intelligence solve the difficult part and then walk away a metre from the finish line. The model gets built, validated, and shown to beat the people it was meant to help — and then remarkably little changes. Forecasts are overridden. Recommendations are ignored. The analyst keeps the old spreadsheet open in the next tab, just in case. A programme funded and defended as a technology investment succeeds on every technical measure and fails on every one that was supposed to justify it.

This essay is about that skipped metre: the human work of adoption. My argument is that we skip it not out of ignorance but out of legibility. The engineering is ownable, schedulable, and demonstrable; the human change is diffuse, politically awkward, and owned by nobody — so it falls off the plan even when every person on that plan agrees it matters. I want to look at what the human side actually consists of, why a training session and a launch email do not touch it, and why the pattern is so stubborn across otherwise very different organisations. I also want to take seriously the strongest objection — that adoption simply follows accuracy, and that change management is a tax the consultancies invented to sell — because it is half right, and the half that is wrong is the more interesting half. These programmes tend to end in the same place: not at the model, but at the organisation staring at how it already handles judgement and authority, and discovering that the technology has done nothing more exotic than hold up a mirror.

2. A metre from the finish line

Picture a demand-forecasting model inside a mid-sized manufacturer. After eight months it is genuinely good: on the held-out data it cuts the forecasting error of the planning team by a comfortable margin, and in the steering meeting the chart that shows this earns a round of applause. It goes live. The planners are given access, a half-day of training, and a cheerful email announcing the new way of working.

Six months later, someone finally looks at the logs. The planners are overriding roughly two in five of the model's numbers — and, this is the part that stings, they are overriding the good ones about as often as the bad ones. Where the model was left alone, error is down by nearly a third against the old baseline. Where it was overridden, the result is marginally worse than the manual process had been before the project started. Net effect on the number the business case promised: close to nothing. The model works. The adoption did not happen.

The specifics stop mattering once you have seen the shape a few times. Swap forecasting for credit decisioning, clinical triage, fraud scoring, or next-best-action in a contact centre, and the story rhymes. Something expensive and clever is built, it clears its technical bar, and then it sits inside the organisation like a graft the body is quietly rejecting. The failure is not in the model. It lives in the metre of ground between a model that works and an organisation that has changed — and that ground is almost never anyone's job.

3. The economics of the skipped step

Ask why the step gets skipped and the usual answer is that people underestimate it. I do not think that is quite right. Most sponsors know, in the abstract, that "the people side" matters — they will say so, unprompted, in the kick-off. The problem is not belief. It is structure.

Consider what the two halves of the work look like from the vantage point of a plan.

Dimension	The technical build	The human change
Owner	A named team, a named lead	Diffuse — "everyone", and therefore no one
Visible progress	Commits, metrics, a demo	Almost nothing to show until late
Definition of done	The model hits its metric	Contested and hard to declare
Political cost	Low — nobody resents a model	High — it touches status and judgement
What cutting it feels like	A failure you must explain	A silence nobody notices

The technical half is legible. It can be owned, scheduled, demonstrated, and defended in a steering meeting with a chart. The human half is illegible in exactly the ways an organisation punishes: it has no clean owner, it produces nothing to show for most of its duration, it cannot easily be declared finished, and — most corrosively — its neglect makes no noise at the moment of neglect. Cutting the model is a visible failure someone will be asked about on the day. Cutting the change work is a silence that only speaks six months later, by which time the programme has been signed off as a success and the team has moved on to the next thing.

This is why the step is skipped even by people who sincerely believe it matters. We are not choosing the model over the people. We are choosing the legible over the illegible, under time pressure, again and again, and the sum of those small and individually rational

choices is a programme that is ninety per cent ready and zero per cent adopted.

The human side of AI adoption is not neglected because leaders think it unimportant. It is neglected because it is illegible – unownable, unshowable, undeclarable – and organisations reliably starve the work they cannot put on a slide.

4. What "the human side" actually is

The phrase invites its own trivialisation. Say "the human side" and most programmes hear "training and communications", budget for a workshop and an intranet page, and consider the box ticked. That is not the human side; it is the packaging. Underneath sit four things, none of which a workshop reaches.

The first is a redistribution of judgement. An AI system does not merely automate a task; it relocates a decision. The planner who used to own the forecast now reviews it. That sounds like a small change and is in fact an enormous one, because the person's expertise has been silently reclassified – from the thing that produces the answer to the thing that second-guesses a machine's answer. Nobody said this out loud, which is precisely why it festers.

The second is trust calibration, and it cuts both ways. Under-trust looks like our forecasting planners: overriding good outputs because they cannot yet tell a good one from a bad one and so default to their own hand. Over-trust is the opposite failure and often the more dangerous – waving through a recommendation in a case where the model's confidence should have been low, because deferring is easier than judging. Calibrated trust is a skill; it is specific to this system's particular strengths and blind spots; and it is learned slowly, through feedback. No system ships with it, and no launch email confers it.

The third is identity. For a great many professionals, the judgement now being mediated by a model is not a task they perform but a large part of who they are at work. The experienced underwriter, the senior planner, the clinician – their standing rests on being the person who knows. Introduce a system that appears to know better and you have not handed them a tool; you have made a quiet argument about their worth. Resistance that presents as "they won't use the thing" is very often, underneath, "you have not told me what I am now for."

The fourth is incentives, which usually go untouched because they sit in another department's gift. If a planner is still measured, rewarded, and blamed exactly as they were before, then every override is rational: they carry the downside of a bad number personally, while the model carries none of it. Asking someone to trust a system that cannot be held accountable, while they remain entirely accountable, is asking them to be brave for no reason. Most people, sensibly, decline.

- Redistribution of judgement – the decision moves, and the expert is reclassified from author to reviewer.
- Trust calibration – knowing when to defer and when to override this specific system, learned slowly and shipped by nobody.

- Identity — for many, the mediated judgement is not a task but a standing, and the system reads as a verdict on their worth.
- Incentives — unchanged accountability makes the human's caution rational and the system's adoption irrational.

None of these four is addressed by explaining how to log in.

5. The objection worth taking seriously

There is a confident counter-argument, and dismissing it would be cheap, so let me put it at full strength. It runs: adoption follows accuracy. Build a system that is unmistakably, repeatedly better, and resistance evaporates on its own, because people are not fools — they adopt what visibly helps them. On this view the whole apparatus of "change management" is a tax invented by consultancies to bill for meetings; the planners in our story were not badly changed, they were correctly sceptical of a model that was not yet good enough, and the remedy is not a workshop but a better model and an honest track record.

The reason to respect this argument is that it is frequently right. A great deal of what gets filed as "resistance" is a rational response to a system that genuinely is not good enough yet, and no quantity of change management should paper over that — nor can it. A tool that quietly saves an individual real effort, every day, with visible reliability, does tend to pull itself into use; and heavy-handed adoption programmes strapped onto weak tools earn the cynicism they attract. Any honest account has to concede that the model's quality is doing a great deal of the work, and that some of what we solemnly call "human change" is really just a system slowly earning the trust it had not yet deserved.

Where the argument fails is in assuming that the individual and the organisation adopt on the same terms. A tool that helps me, privately, spreads on its own. But most enterprise AI does not help the individual privately; it reallocates judgement and risk across people. It helps the organisation precisely by changing what the planner, the underwriter, the clinician is for, and it asks each of them to carry a personal risk on behalf of a collective gain they may never directly see. "Accuracy sells itself" is true for the spreadsheet on my own laptop and false for the system that quietly rewrites my role and leaves my incentives pointing the other way. The objection is an accurate description of consumer software adoption and a misleading one for institutional change. It is right about the tool and wrong about the organisation.

6. Speed, trust, and the tension no method dissolves

There is a genuine tension here that no method dissolves, and an essay should sit with it rather than pretend to resolve it. The whole promise of AI is speed — faster decisions, more of them, at lower cost. But trust, of the calibrated kind described above, is built slowly, through exactly the feedback loops that speed is designed to eliminate. The faster a system decides, the fewer chances a human has to learn where it can be believed. We are, in effect, asking people to extend more trust, faster, to a thing that is harder to check because it is fast. Push adoption too hard and you manufacture over-trust — a workforce

nodding through recommendations it no longer understands. Push too gently and the system never escapes the parallel-spreadsheet purgatory.

The organisations that navigate this well tend to do something that looks, from outside, like hesitation. They keep the human firmly in the loop long past the point where the metrics say the model could safely run alone — not because the model still needs the human, but because the trust does. Calibrated confidence cannot be back-filled after go-live; it has to be laid down during it, case by case, while there is still a person watching who can be surprised. This is expensive and it reads as timidity on a status report, and it is the single most reliable marker I know of a programme that will still be delivering value in two years rather than two quarters.

“Adopting AI is a change to people's work — and changes to people's work obey the same unglamorous rules they always have.”

7. What taking it seriously looks like

None of this points toward a heavier change-management bureaucracy; more workshops would only add cost to the illegible column of the ledger. It points toward a few shifts in how the work is framed and owned.

1. Fund the change as part of the build, not as a wrapper around it. If the human work has no owner, no visible milestones, and no definition of done, it will be starved regardless of anyone's good intentions. Give it a name, a person, and something concrete to show at each checkpoint.
2. Design the human's new job on purpose. If you are moving someone from author to reviewer, say so, out loud, and define what good reviewing now is — when to defer, when to override, what they are uniquely responsible for that the model is not. An unspoken reclassification is the thing people resist; a deliberate, dignified one is something they can step into.
3. Fix the incentives before you blame the behaviour. If overriding is rational given how someone is measured, then change the measure or expect the override. You cannot exhort your way past an incentive that points the other way.
4. Treat calibrated trust as a deliverable, not a by-product. Show people where the system is strong and where it is blind, in their own cases, with feedback, over time. Trust is a skill you can teach — but only slowly, and only with honesty about the model's limits.

There is nothing proprietary in any of this. It is the recognition that adopting AI is a change to people's work, and that changes to people's work have always obeyed rules that predate the technology by a century.

8. The mirror

The deepest reason the human side gets skipped may be that it is not really about AI at all, and organisations would much rather it were. A technology problem can be bought, scoped, and delegated to a team with a metric. The problems that AI adoption actually surfaces — who owns a decision, whose judgement counts, how honestly the place handles the demotion of expertise, whether its incentives point where its strategy claims to — are old, political, and unowned. They were there before the model and they will be there after it. The model did not create them; it made them impossible to ignore, briefly, before the organisation found a way to ignore them again.

That is the pattern worth naming. An AI programme is a remarkably efficient instrument for revealing how an organisation already treats change and authority. Where a place is honest about power and generous about identity, adoption does tend to follow the model's quality more or less as the optimists promise. Where it is not, no accuracy figure will save the programme, because the resistance was never really about the model. We keep funding the metre we can see and skipping the metre we cannot, and then calling the result a technology failure — when the technology, for once, did exactly what it said it would.

The work everyone skips is not a phase of the project. It is the project. We have simply been calling it by the name of the easier half.