

BEYOND THE ALGORITHM

Change Management for AI Adoption — The Human Side Everyone Skips

A Practitioner's Perspective on Why Organisations Invest in Algorithms but
Not in the People Who Must Use Them

Giovanni Leonardi

September 2022



Change Management for AI Adoption — The Human Side Everyone Skips

Transformation · No. 287

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

“The organisations that succeed with AI are not the ones with the best models — they are the ones that took the time to understand what would break when the model arrived.”

1. The Investment Asymmetry

There is a pattern that recurs with striking consistency across organisations adopting artificial intelligence. The technology budget is approved. The data science team is recruited or contracted. The platform is selected, the models are built, and the pilot delivers promising results. Then the organisation attempts to move from pilot to production — and discovers that the hardest part of AI adoption has nothing to do with AI.

The hardest part is the people.

This observation should not be surprising. Every significant technology shift in the last three decades — enterprise resource planning, customer relationship management, digital channels, cloud migration — has eventually confronted the same truth: that technology changes are, at their core, human changes. Yet with AI, the pattern is repeating itself with a fidelity that borders on wilful. Organisations are investing heavily in the technical infrastructure of AI while treating the human infrastructure as an afterthought, if they treat it at all.

The result is a growing inventory of technically sound AI capabilities that nobody uses, that people actively work around, or that deliver a fraction of their potential because the organisation around them was never prepared to absorb them.

2. Why the Human Side Gets Skipped

The reasons are structural, not accidental, and understanding them matters more than simply lamenting the gap.

AI programmes are overwhelmingly technology-led. In most organisations, the AI agenda sits with the Chief Technology Officer, the Chief Data Officer, or a newly created AI

function that reports into technology. The programme teams are dominated by data scientists, machine learning engineers, and platform architects — people who are exceptionally good at building models and exceptionally unlikely to spend their time thinking about how a claims handler in Birmingham will react when their workflow changes. This is not a criticism of those teams. It is a structural observation about where the centre of gravity sits and what it optimises for.

Vendor and consultancy incentives reinforce the technology bias. The organisations that sell AI capability — whether cloud platforms, specialist tooling, or advisory services — sell technology. Their commercial models are built around platform licensing, model development, and technical implementation. Change management, if it appears at all, is a line item in the statement of work that is the first to be descope when budgets tighten. The market, in other words, is structured to oversupply the technical dimension and undersupply the human one.

Change management teams are not equipped for AI. This is perhaps the most uncomfortable truth. The change management profession has developed sophisticated methods for managing technology-driven change — stakeholder analysis, communications planning, training design, adoption measurement. These methods work well for changes that are predictable, well-bounded, and largely procedural: a new system replaces an old system, and the change programme helps people make the transition. AI does not behave like this. A machine learning model does not simply replace a process step; it changes the nature of judgement, the distribution of authority, and the meaning of expertise in ways that conventional change methods are not designed to address.

The human side is harder to quantify. Technology investment produces tangible, measurable outputs — a deployed model, an API endpoint, a prediction accuracy score. The human side produces outcomes that are diffuse, delayed, and difficult to attribute: trust, adoption, confidence, willingness to change working practices. In organisations where investment cases are built on measurable returns, the things that are easy to measure receive the money, and the things that matter most often do not.

3. What Change Management for AI Actually Requires

If conventional change management is insufficient, then what does a more adequate approach look like? In my experience, it requires organisations to engage with at least four dimensions that most current programmes neglect.

The trust dimension. AI adoption is fundamentally a trust problem, and trust operates differently from competence. A person can be perfectly capable of using an AI-powered tool and still refuse to rely on it — not because they lack training but because they do not trust the output. This is rational behaviour, not resistance. Building trust requires transparency about what the model does and does not do, sustained exposure to the model's behaviour in low-stakes settings, and — critically — permission to override the model when human judgement says otherwise. Organisations that deploy AI as a fait accompli, presenting model outputs as decisions rather than recommendations, systematically undermine the trust they need.

The identity dimension. This is the dimension that receives the least attention and may matter the most. When an AI model takes over a task that a person has performed for years — assessing credit risk, triaging customer enquiries, reviewing documents for compliance issues — it does not simply change what that person does. It changes who they are, professionally. The underwriter whose expertise lay in assessing risk now supervises a model that assesses risk. The compliance analyst whose value was their judgement now reviews the judgement of an algorithm. These are not trivial transitions. They go to the heart of professional identity, and they provoke responses — anxiety, defensiveness, disengagement — that no amount of training can address if the identity question is not confronted directly.

The workflow dimension. AI rarely arrives as a clean substitution. More often, it reshapes workflows in ways that create new friction, new handoff points, and new ambiguities about who is responsible for what. A model that automates the first stage of a process may speed up that stage while creating a bottleneck downstream, where human reviewers now face a higher volume of pre-processed cases that still require judgement. The workflow around the AI matters as much as the AI itself, and designing that workflow is a change management task, not a technology task.

The governance dimension. When a model makes or influences a decision, who is accountable? This question sounds straightforward until an organisation tries to answer it in practice. The data scientist who built the model? The business owner who approved its deployment? The operator who acted on its recommendation? In most organisations, the accountability framework for AI-assisted decisions is either absent or so vague as to be meaningless. This is not merely a compliance risk — it is a change management problem, because people will not adopt a tool if they do not understand who carries the responsibility when it gets things wrong.

4. The Change Function Itself Is Not Ready

There is a deeper issue that the profession has been slow to confront. Change management, as it is practised in most organisations, was built for a world of deterministic technology — systems that do what they are designed to do, every time, in the same way. AI is probabilistic. It produces different outputs from similar inputs. It degrades over time as data distributions shift. It exhibits behaviours that its own creators cannot fully explain. Managing the human response to a technology that behaves like this requires a different kind of change practice — one that is comfortable with ambiguity, equipped to facilitate ongoing adaptation rather than one-time transition, and honest about what it does not know.

The pattern I have observed across sectors is that change teams, when confronted with AI, default to what they know. They produce stakeholder maps, communications plans, and training curricula — the standard toolkit — and apply it to a challenge that the toolkit was not designed for. The result is change activity that is visible, well-structured, and largely ineffective.

This is not a failure of effort or intent. It is a failure of fit.

5. A More Honest Starting Point

The organisations that succeed with AI are not the ones with the best models — they are the ones that took the time to understand what would break when the model arrived. This understanding does not come from a stakeholder analysis conducted by a project team in isolation. It comes from a genuine, sustained engagement with the people whose work will change — an engagement that starts before the model is built, not after it is deployed.

The question is not "how do we get people to adopt the AI?" The question is "how does work need to change, and what do people need in order to change with it?"

This reframing matters because it shifts the burden. The first question puts the problem on the people — they need to adopt. The second question puts the problem on the organisation — it needs to create the conditions for change. In my experience, organisations that ask the second question consistently outperform those that ask the first, because they invest in the things that actually determine whether AI delivers value: redesigned workflows, clarified accountability, rebuilt professional identities, and earned trust.

None of this is easy. None of it is fast. And none of it fits neatly into a technology programme plan with fixed milestones and a go-live date. But the alternative — continuing to invest in the algorithms while ignoring the people — is a strategy that has already failed, in organisation after organisation, across every sector I have worked in.

The human side is not a phase to be managed. It is the work itself.