


BEYOND AUTOMATION ANXIETY

Delivery in 2030 — Four Scenarios for the Profession That Builds Change

Two Uncertainties, Four Futures, and the Capabilities Worth Building Now

Giovanni Leonardi

April 2026



Contents

1. Executive Summary 3

2. The Scene That Keeps Recurring 3

3. Two Axes of Uncertainty 4

 3.1 How Far Does Agent Autonomy Actually Get? 4

 3.2 Do Organisations Restructure — or Bolt It On? 4

4. The Scenario Matrix 5

 4.1 Augmented Craft 5

 4.2 The Thin Orchestra 6

 4.3 Bolt-On Purgatory 6

 4.4 Trust Recoil 7

5. No-Regret Moves 8

 5.1 Judgement Under Ambiguity 8

 5.2 Agent Fluency 8

 5.3 Governance That Actually Governs 8

6. What We Cannot Know 9

Delivery in 2030 — Four Scenarios for the Profession That Builds Change

Programme · No. 253

April 2026

“The conductor who sets the wrong tempo does not merely slow the performance — she misdirects fifty agents simultaneously, compounding errors at machine speed.”

1. Executive Summary

The delivery profession — programme managers, portfolio leads, transformation directors — faces a period of genuine uncertainty. AI agents are absorbing tasks that once defined the role: status consolidation, dependency tracking, risk pattern recognition, elements of stakeholder communication. The question is no longer whether agents will reshape delivery work, but how far that reshaping extends and whether organisations will redesign themselves around it or simply layer it over what already exists.

This paper applies scenario method to that question. It identifies two axes of honest uncertainty — the degree of agent autonomy that actually materialises in enterprise delivery, and the willingness of organisations to restructure around it — and explores the four futures they produce: Augmented Craft, The Thin Orchestra, Bolt-On Purgatory, and Trust Recoil. None is a prediction. Each describes a world with its own internal logic, its own winners and losers, and its own implications for how we practise. The strategic value of the exercise lies in what emerges at the intersection: three capabilities — judgement under ambiguity, agent fluency, and governance that actually governs — that pay off in every scenario. Those are the no-regret moves, and they are worth building now.

2. The Scene That Keeps Recurring

The specifics vary, but the shape is always the same. A programme director walks into a Monday stand-up to find that an automated dependency analysis, run overnight, has not only identified a critical-path conflict between two workstreams but has drafted three resolution options, modelled the schedule impact of each, and pre-populated a decision paper for the steering committee. Work that would have taken a senior planner two days has been completed in forty minutes.

The reaction is never quite celebration. It is something quieter and more unsettling — a professional recognising that the mechanical substrate of her expertise is being absorbed, and not yet knowing what sits underneath it.

We are all living some version of this scene in early 2026. The agentic systems entering enterprise environments are not the chatbots of three years ago. They orchestrate multi-step workflows, reason across structured data, and produce outputs that require correction rather than creation. Status reporting, risk consolidation, schedule optimisation, first-pass stakeholder analysis — these are not the future of agent capability in delivery. They are the present, deployed in production across a growing number of organisations.

But the gap between what agents can do in a controlled demonstration and what they reliably accomplish inside a complex, politically charged transformation programme remains substantial. And the gap between what technology makes possible and what organisations actually choose to reorganise around is, historically, wider still. These two gaps — in capability and in will — are the uncertainties that will shape our profession through the rest of this decade.

3. Two Axes of Uncertainty

Scenario planning is not forecasting. Its value lies in resisting the temptation to call the outcome and instead mapping the territory of plausible futures. The method requires identifying the uncertainties whose resolution genuinely determines which world we inhabit — then holding all possibilities open simultaneously.

For delivery work through 2030, two uncertainties dominate.

3.1 How Far Does Agent Autonomy Actually Get?

The first uncertainty is technical in origin but organisational in expression. How far does autonomous agent capability extend in the messy, ambiguous, politically textured environment of enterprise delivery?

The optimistic trajectory is visible: agents move from task execution to judgement, from structured to unstructured work, from single-domain to cross-domain orchestration. On this path, by 2028 or 2029, an agent fleet manages the operational rhythm of a programme — tracking commitments, flagging deviations, adjusting plans, escalating only the genuinely novel — while human oversight concentrates at the strategic edges.

The sceptical trajectory is equally credible. Agent capability plateaus at the boundary of structured work. The messy middle of delivery — the ambiguous stakeholder signal, the political subtext in a steering paper, the judgement call about when to push a supplier and when to absorb a delay — remains stubbornly human. Agents become superb at mechanics and unreliable at everything else. Autonomy stalls not because the technology fails outright, but because the operating environment is too noisy, too political, and too context-dependent for agents to navigate without constant human correction.

Both trajectories are consistent with what we observe today. The honest answer is that we do not know which will hold.

3.2 Do Organisations Restructure — or Bolt It On?

The second uncertainty is entirely human. When agent capability arrives — at whatever level — do organisations redesign their delivery structures around it, or do they layer it over what they already have?

Restructuring means genuinely rethinking team composition, role definitions, reporting lines, and governance rhythms. It means acknowledging that a programme team of forty, with powerful agent support, might deliver the same outcome as a team of twelve with different skills. It means confronting the career models, the organisational charts, and the consulting economics that depend on headcount.

Bolting on means doing what large organisations almost always do with new technology: layering it over existing structures, keeping the same team shapes, the same governance cadences, the same role definitions, and expecting the technology to make everything a bit faster and a bit cheaper. The org chart stays untouched. The agents fill gaps rather than replace functions.

History favours bolt-on, heavily. Organisations are structurally conservative; they absorb new capability into existing shapes rather than redesigning around it. But the economic pressure this time may differ — the potential headcount leverage of autonomous agents creates a cost argument that is difficult to set aside in a sustained cost-management environment. Whether that pressure is enough to overcome institutional inertia is precisely the uncertainty.

4. The Scenario Matrix

These two uncertainties create four distinct worlds for delivery work. Each is internally coherent, each carries different implications for the profession, and none is more likely to be right than any other. The point of the exercise is to inhabit all four simultaneously.

	Organisations Restructure	Organisations Bolt On
High Agent Autonomy	The Thin Orchestra	Trust Recoil
Limited Agent Autonomy	Augmented Craft	Bolt-On Purgatory

4.1 Augmented Craft

Limited autonomy · Deep restructuring

Agents prove excellent at structured mechanics — scheduling, dependency tracking, status consolidation, test automation, document generation — but hit a hard ceiling at judgement, ambiguity, and political navigation. Organisations, however, take the restructuring seriously. They slim delivery teams, redesign roles around the new human-agent boundary, and invest in upskilling the practitioners who remain.

The result is a genuine elevation of the craft. Delivery professionals spend less time on mechanics and more on the work that agents cannot reach: navigating organisational politics, building coalitions for difficult decisions, translating between technical reality and executive expectation, exercising the contextual judgement that emerges from years

of pattern recognition in complex environments.

Seniority matters most in this world. Junior delivery roles contract sharply — the entry-level tasks that once trained new practitioners are now handled by agents. The profession becomes harder to enter but more rewarding to practise. The programme manager becomes something closer to a strategist-diplomat than a coordinator, and the skills that define the role shift decisively toward political acuity, systems thinking, and the capacity to make sound decisions under genuine ambiguity.

The risk is a talent pipeline that quietly dries up. If the apprenticeship ladder disappears — if there is no longer a graduated path from junior planner to senior programme director that teaches judgement through accumulated responsibility — the profession may elevate its current generation while failing to develop the next one.

4.2 The Thin Orchestra

High autonomy · Deep restructuring

Agent capability crosses the judgement threshold. Agents handle not only structured mechanics but substantial elements of unstructured work: stakeholder sentiment analysis, risk interpretation, adaptive replanning, preliminary negotiation of scope trade-offs. And organisations restructure to exploit it fully.

The result is radical compression. A programme that once required a team of forty — planners, coordinators, PMO analysts, test managers, change leads — now runs on a core of five or six humans conducting a fleet of specialised agents. The metaphor is orchestral: a small number of skilled conductors directing a large ensemble of capable performers, with value concentrated in the human ability to set direction, interpret context, and make the calls that carry genuine accountability.

This is the highest-leverage scenario for those who remain. The delivery leader in this world commands resources that would have been unimaginable five years earlier. But the economics are brutal for the profession at large. Delivery headcount collapses. The consulting firms that built their revenue models on placing programme teams face an existential challenge. The large PMOs that defined delivery governance in complex organisations are hollowed out.

The counterweight is that the thin-orchestra model demands an extraordinarily high calibre of human practitioner. The conductor who sets the wrong tempo does not merely slow the performance — she misdirects fifty agents simultaneously, compounding errors at machine speed. The premium on judgement, experience, and contextual intelligence is the highest of all four scenarios.

4.3 Bolt-On Purgatory

Limited autonomy · Shallow restructuring

This is the friction scenario — and, if we are candid, probably the nearest-term default for most large organisations. Agent capability stalls at structured work. Organisations, true to form, layer it over existing structures without redesigning anything.

The result is purgatory: agent tools proliferate, each team adopts different ones, integration is poor, governance is improvised, and the net effect on delivery productivity is far smaller than anyone promised. Programme managers spend as much time managing the agents — correcting outputs, arbitrating between conflicting automated recommendations, explaining to stakeholders why the AI-generated status report contradicts what they observe on the ground — as they save by using them.

The delivery profession survives largely intact in this world but grows measurably more frustrated. The promise of agent-assisted delivery collides with the reality of fragmented tooling, patchy data quality, and organisational resistance to changing the processes that agents are supposed to improve. The most common experience is the agent that produces a technically correct analysis from technically correct data that is nonetheless wrong, because it lacks the contextual knowledge — the unwritten dependencies, the political agreements, the historical commitments — that a human programme manager carries.

Bolt-On Purgatory is also the highest-friction scenario for agent vendors. The gap between what their products achieve in a clean demonstration and what they accomplish inside a real enterprise — with real data quality, real legacy architecture, and real organisational politics — becomes their central commercial problem. The disillusionment cycle that followed every enterprise technology wave of the past two decades plays out once more.

4.4 Trust Recoil

High autonomy · Shallow restructuring

This is the scenario that ought to concern governance practitioners most. Agent capability reaches genuinely high levels — agents can operate autonomously across substantial portions of a delivery programme — but the organisational structures, governance frameworks, and regulatory environment have not adapted. High-autonomy technology runs inside low-maturity governance, until something goes visibly, publicly wrong.

The triggering event need not be catastrophic; it must be public and attributable. An agent fleet in a major infrastructure programme makes autonomous scheduling decisions that cascade into a safety-relevant error. A financial transformation discovers that commitments have been made — to regulators, to partners — that no human reviewed or approved. The political and regulatory response is swift: mandatory human-in-the-loop requirements, agent decision audits, new liability frameworks, possibly sector-specific restrictions on autonomous delivery systems.

The profession in this scenario experiences a recoil toward human accountability. The delivery leader becomes, above all, the person accountable for what the agents do — the role defined not by what it produces but by what it signs off. Governance layers expand. Assurance functions multiply. The efficiency gains of high autonomy are partially consumed by the compliance overhead required to maintain public and regulatory trust.

Trust Recoil is not necessarily a poor outcome for the profession — it preserves the centrality of human judgement and creates new roles in agent governance and assurance. But it is an inefficient one: the technology can do more than it is permitted to do, and the gap between capability and permission becomes a persistent source of organisational

tension.

5. No-Regret Moves

The strategic value of scenario planning lies not in choosing between futures but in identifying the actions that pay off regardless of which future arrives. Across all four scenarios, three capabilities emerge as robust — the investments that are never wasted.

5.1 Judgement Under Ambiguity

In every scenario — whether agents handle mechanics alone or cross into substantive work — the irreducible core of human value in delivery is the ability to make sound decisions when information is incomplete, the politics are complex, and the stakes are high. This is not a capability that most delivery training develops deliberately. We train people in method, in tooling, in process. We rarely invest in the disciplined exercise of judgement — the ability to weigh competing signals, to recognise when to push and when to yield, to make the call that cannot be delegated to a system however capable.

The capabilities worth building now are those that pay off in all four futures. Judgement under ambiguity sits at the top of that list — not because it is new, but because every scenario makes it more consequential and nothing else substitutes for it.

Building this capability means investing in apprenticeship structures that expose practitioners to ambiguity at graduated levels of consequence, in reflective practice that turns experience into transferable pattern recognition, and in the kind of honest feedback cultures where poor judgement calls are examined without blame and good ones are understood rather than merely celebrated. None of this is novel. What changes is the urgency.

5.2 Agent Fluency

Not AI expertise. Not prompt engineering. Agent fluency: the practical, professional-grade ability to work alongside, direct, supervise, and correct autonomous systems in a delivery context. This means understanding what agents are actually doing — not only their outputs but their reasoning and their failure modes. It means knowing how to calibrate trust in agent-generated analysis, recognising the characteristic patterns of confident error that current agentic systems produce, and being able to intervene effectively when they go wrong.

Agent fluency is not a specialism to be siloed in a technical team. It is a baseline professional competence, as fundamental to the next decade of delivery work as spreadsheet fluency was to the last two. Every delivery role — from junior coordinator to portfolio director — will need it, and the profession's training infrastructure has not yet begun to build it at the scale required.

5.3 Governance That Actually Governs

We have spent years building governance frameworks, and most of them govern reporting rather than decision-making. They track status upward without enabling decisions downward. This weakness is chronic in every scenario, but it becomes acute in a world of agent-assisted or agent-led delivery, where the speed and volume of operational decisions outpaces any governance model built around monthly steering committees and quarterly stage gates.

The governance model that works across all four scenarios defines decision rights clearly, delegates authority at the level where information lives, and creates fast, auditable escalation paths for decisions that exceed delegated authority. This is not new thinking — it is a principle that most governance frameworks articulate and few actually implement. The arrival of agents makes the gap between principle and practice untenable. In the Trust Recoil scenario, it is the gap that causes the crisis. In The Thin Orchestra, it is the gap that determines whether the model works or collapses under its own leverage. In every scenario, closing it is the single most important structural preparation the profession can make.

6. What We Cannot Know

We are in the early months of a shift whose contours are not yet settled. The future will almost certainly contain elements of more than one scenario: restructuring may proceed unevenly across sectors, with regulated industries operating closer to Trust Recoil while technology firms move toward The Thin Orchestra. The autonomy boundary may shift differently across domains of delivery work, reaching deep into scheduling and risk analysis while remaining shallow in stakeholder management and benefits realisation. The triggering events that would push one scenario into dominance may not have happened yet — or may already be unfolding in ways we have not recognised.

“We have spent our careers helping organisations through change. The question now is whether we can do the same for ourselves.”

What we can know is that the profession which builds change is itself facing a change it cannot manage with its existing tools. The methods, career structures, and organisational models that served delivery work for the past two decades were designed for a world in which the fundamental unit of delivery capability was the human practitioner. That assumption is now, at minimum, under serious question. The practitioners who navigate this well will be those who approach it as they would any complex transformation: with clear-eyed assessment of what is actually happening, disciplined preparation for multiple outcomes, and the professional confidence to adapt their own practice to a world they did not design and cannot fully predict.