


BEYOND THE GREEN DASHBOARD

The Part of the Job the Agents Cannot Take: Programme Management in an Age of Autonomous Delivery

Why the parts of the role that resist automation are the parts that were always the job — a view from inside the programme

Giovanni Leonardi

April 2024



The Part of the Job the Agents Cannot Take: Programme Management in an Age of Autonomous Delivery

Programme · No. 258

Originally written April 2024 | Re-edited and re-rendered for the WhereBeyond Edition, 2025

"The hollowing-out will not happen because the machines can do the whole job. It will happen because we never made the case for the part they can't."

1. The dashboard is green and nothing is moving

The dashboard is green. Every workstream reports on plan, the milestone view shows amber at worst, and the summary slide carries all the reassuring vocabulary of control: mitigations in place, risks actively managed, delivery confidence high. And yet anyone who has sat in enough of these rooms can feel the thing the pack is built to hide. The programme is not moving. It has not moved for six weeks. Beneath the green there is a decision nobody will own — who funds the shared platform, whose operating model gives way — and the reporting has quietly organised itself around not naming it.

I begin here because this scene is at once the most threatened and the least threatened by the arrival of autonomous software agents into delivery work. The current wave of generative tools can already assemble that status pack. Point them at the plan, the risk log and the last month of updates and they will produce a cleaner, faster, more internally consistent report than the two analysts who spend three days a fortnight doing it by hand. What they cannot do is the only thing that would have helped: walk into the room and make the two people who must agree actually agree.

That gap — between the work the machines are about to take and the work that was always the point — is the whole question for our profession right now. And we are answering it badly, because we have spent twenty years describing programme management as the very thing the agents do best.

2. What the agents genuinely take

Let me concede the disruptive case in full, because the honest version of it is stronger than the anxious one.

A great deal of what fills a programme manager's week is coordination overhead: chasing updates, reconciling versions of the plan, formatting the pack, translating one team's status into another team's language, maintaining the RAID log, preparing the pre-read,

writing the minutes, tracking the actions. This is real work and it has real value — an unmaintained plan is a dangerous plan — but almost none of it requires the judgement we claim as our craft. It requires diligence, memory and time. Those are precisely the qualities software is good at.

- An agent can hold nine hundred dependencies in view and flag the three that just became critical, without the fatigue that makes a human miss the fourth.
- It can reconcile six workstream trackers into one truth overnight, and tell you where they disagree rather than papering over it.
- It can draft the status narrative, the risk summary and the steering pre-read, and do it in the house style, in minutes.
- It can watch the plan continuously rather than in fortnightly snapshots, so that slippage surfaces the day it happens rather than the day before the board.

Consider a composite that will be familiar to anyone who has run a large change portfolio. A programme of roughly forty people, a fourteen-million-pound budget, a status pack that consumes two analysts for the better part of three days every fortnight. Give the reporting to a competent agent and that cost very nearly disappears — call it a hundred-odd analyst-days a year handed back to the programme. That is not a trivial saving, and it is coming whether we welcome it or not.

If your honest answer to what do you do all week is a list of things an agent could do overnight, the anxiety you feel about these tools is not misplaced. It is information.

So the disruptive case lands. The reporting layer, the coordination layer, the administrative scaffolding of delivery — much of it is about to become cheap. Anyone whose role is mostly that scaffolding should take the warning seriously.

3. The steelman — and where it breaks

The strongest version of the obsolescence argument goes one step further, and it deserves a real answer rather than a reassuring one. It says: coordination is not a layer on top of the job; coordination is the job. Strip out the chasing and reconciling and reporting and there is no residual role of any size left — just a thin band of leadership that the sponsor was doing anyway. On this view the programme manager is middle-management connective tissue, and agents dissolve connective tissue for a living. Flatten the layer, wire the teams directly to the agents, and let the accountable executive make the calls.

I take this seriously because in a certain kind of programme it is correct. Where the work is genuinely modular, the interfaces are clean, the funding is settled and the political questions have already been answered, a heavy coordination layer is waste, and it always was. Agents will expose that waste ruthlessly, and some programme management roles will disappear — as they should.

But the argument smuggles in an assumption that the interesting programmes never satisfy: that the hard part of delivery is keeping track of the work. It isn't. The hard part is the decisions the work keeps running into — the ones that are ambiguous, contested, and owned by no single person. Coordination is not connective tissue in those programmes. It

is the visible residue of a continuous, mostly invisible act: getting people who are not incentivised to agree to commit anyway. The reporting is downstream of that. Automate the reporting and the act itself does not shrink; it is simply left standing in the open, without the busywork that used to disguise how much of the job it really was.

“An agent can tell you the truth about your plan far faster than a frightened team will. It cannot make the two people who must agree actually agree.”

4. What remains when the scaffolding is gone

Take away the administrative accretions and three things are left exposed. None of them is new. All of them were always the job.

The first is judgement under ambiguity. A plan tells you a milestone is at risk; it does not tell you whether to hold the date and absorb the pain, move it and spend the credibility, or descope and fight the benefits argument later. That call depends on things not in any dataset the agent can see: how much political capital the sponsor has left, whether the regulator is in a patient mood this quarter, which of your suppliers is bluffing. An agent will happily generate three options and even a recommendation. Someone still has to own the consequence of choosing, in a room full of people who will remember the choice.

The second is the brokering of commitments. Most of what actually unblocks a stuck programme happens in the ten minutes before the meeting and the corridor afterwards — the quiet trade where operations agrees to carry a risk if finance moves a date, sealed with nothing more binding than two people's word and reputation. This is not documented because it cannot be; the moment you try to formalise it, it stops working. An agent has no reputation to stake and no corridor to stand in. It can record the commitment once made. It cannot make it.

The third is political navigation, which is the least respectable phrase in our vocabulary and the most important skill in the work. Every significant programme is a negotiated settlement between parts of an organisation that want different things. Reading that field — who has really conceded and who is waiting to relitigate, when a decision is safe to force and when forcing it will detonate — is the discipline that separates programmes that land from programmes that are technically well-run and quietly fail. It is judgement about people, and it is exactly the register in which today's tools are weakest and most confidently wrong.

1. The agent surfaces the decision. The programme manager owns it.
2. The agent records the commitment. The programme manager wins it.
3. The agent maps the stakeholders. The programme manager reads the room.

5. The real danger is not obsolescence

Here is the uncomfortable turn, and it is aimed at us rather than at the technology. The profession is in danger — but not of being replaced. It is in danger of being misunderstood at exactly the wrong moment.

For two decades we have described ourselves in the language of the automatable part. We certify people in the mechanics of planning and reporting. We hire against tool proficiency and RAID discipline. We promote the person with the tidiest pack, not the one who quietly made the intractable decision happen. We have, in short, priced the scaffolding and given the real work away for free — precisely because it is hard to see, hard to measure, and awkward to name in a business case.

Now the scaffolding is about to be commoditised. If an organisation believes the scaffolding was the job — and we have spent years telling it exactly that — then the rational move is to buy the agents and cut the people. The hollowing-out will not happen because the machines can do the whole job. It will happen because we never made the case for the part they can't.

The threat is not that agents will do what programme managers do. It is that organisations will believe they can — because we described the job as the part that automates, and stayed silent about the part that doesn't.

6. Letting the machines take the labour

The response is not to resist the agents. Resisting them is both futile and foolish, because they genuinely make delivery better. The response is to let them take the scaffolding, and to spend the returned time on the work that was always underpriced.

Concretely, that means three shifts. Move from producing reports to interrogating them — if the agent writes the status, the programme manager's job is to know which green is real. Move from tracking commitments to manufacturing them — treat the securing of a contested decision as the primary deliverable, not the paperwork around it. And move, in how we describe and develop the role, from the language of administration to the language of judgement — because the version of programme management that survives the next few years is the one that can say out loud what it is for.

The dashboard will still be green. But someone in the room has to be able to look at it and say: this is green because the plan is healthy, or this is green because we are all frightened of the same decision. The agent will not tell you which. That someone — that judgement — was always the job. The machines are simply about to make it impossible to pretend otherwise.