


BEYOND THE ARTEFACT

AI Made the Programme Report Nearly Free — and That Is the Danger

The tools automated the artefacts of programme delivery; the judgement and accountability they were mistaken for did not move an inch

Giovanni Leonardi

June 2024



AI Made the Programme Report Nearly Free — and That Is the Danger

Programme · No. 225

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

“Accountability is the human willingness to carry consequences, and a system that carries no consequences cannot hold it.”

The report that had never looked better

The status report arrived the night before the steering committee, and it was, by some distance, the best-looking one the programme had produced. The narrative was crisp. The risks were phrased with a confidence the previous versions had lacked. The milestone commentary read as though written by someone who had slept. It had, in fact, been assembled in under a minute by an assistant pointed at the programme's own trackers, and lightly tidied by a delivery lead who had better things to do at nine in the evening. It was also, in the one way that counted, wrong: it reported the integration workstream as green when anyone who had sat in that week's technical calls knew it was amber sliding toward red. The data the assistant had read said green. The data was three days stale, and it had never captured the thing that was actually going wrong.

This is the question the programme profession has been asking itself all year, usually in a slightly anxious register: what does AI do to project and programme management? The honest answer is more interesting than either the marketing or the anxiety suggests, and it comes in two halves that are usually run together. A great deal of what programme managers produce has just become nearly free to produce. Almost none of what programme management is has changed at all. Confusing those two is the mistake I have watched capable people make repeatedly this year, and it is more dangerous than simply being sceptical of the tools.

What the tools genuinely change

Begin with the real gains, because they are real and the sceptics undersell them. The production of programme artefacts — the status reports, the first-cut plans, the RAID entries, the meeting minutes, the endless reformatting of the same facts for different audiences — has been the invisible tax on the profession for as long as it has existed. That tax has genuinely fallen. An assistant will now turn a transcript into actions in seconds, draft a stakeholder update in three registers, and interrogate a plan for missing dependencies faster than a PMO analyst working by hand. Used well, this hands time back to the people who were spending their evenings on formatting rather than thinking.

There is a subtler gain too, and it is the one worth protecting. A model that can read across every workstream's notes at once will surface patterns a single overstretched programme manager might miss: the same risk described in three different workstreams under three different names, the dependency that everyone assumed someone else owned. As a second pair of eyes over a large and messy information estate, the tooling earns its place. This is the part I would defend against the reflexive sceptics — the assistance is not cosmetic, and a programme manager who refuses it on principle is choosing to be slower for no gain.

What they leave exactly where it was

Now the other half, which the enthusiasts skate over. Programme management, stripped to its core, is not the production of artefacts. It is judgement exercised under uncertainty and accountability that cannot be delegated. Neither has moved.

The judgement is the work of deciding which amber is a genuine amber and which is a nervous project manager over-reporting; which slipping milestone is noise and which is the first tremor of a structural problem; which stakeholder's stated objection is the real one and which is a proxy for a concern they will not say aloud. Models are trained to produce the fluent and the plausible, and a stalling programme very often looks, on paper, exactly like a healthy one — right up until it does not. The assistant that wrote the beautiful green report was not malfunctioning. It was doing precisely what it is built to do, rendering the available data into a confident narrative, and the available data did not contain the problem.

The tools have made it effortless to produce an artefact that looks like control. They have done nothing to produce control. Those were never the same thing — and this year the gap between them became easier to hide.

Then there is accountability, which is even more stubborn. When a programme fails, someone stands in front of a board and owns it. That role cannot be delegated to a model — not for reasons of technical maturity but for reasons of meaning: accountability is the human willingness to carry consequences, and a system that carries no consequences cannot hold it. Everything upstream of that moment can be assisted. The moment itself cannot.

The comfortable error

Here is where the two halves collide, and where the profession's real risk lies this year. Programme management has always been able to half-hide behind its artefacts. A thick, well-formatted report has often stood in for genuine control, and an experienced steering committee learns to read past the document to the person presenting it, watching whether the commentary is lived-in or merely typed. What the tools have changed is the cost of the disguise. When a convincing report took a day to write, the effort was itself a weak signal that someone had at least engaged with the material. Now

that a convincing report takes a minute, that signal is gone. The artefact and the work have been decoupled, and the artefact has never looked more reassuring.

The organisations getting this right this year are the ones that quietly shifted what they trust. They stopped treating the polished status report as evidence of anything and started treating the conversation around it as the actual instrument: the fifteen minutes of unscripted questions that no assistant can generate, because the answers do not exist in any tracker. Ironically, the arrival of AI has made the oldest and most human part of programme assurance more important, not less. When the document can no longer be trusted to reflect effort, the only remaining test is the judgement of the person in the room.

What this means for the profession

So the anxious question — will AI replace the programme manager? — has the shape of the answer built into its own error. It assumes the job was the production of the artefacts, which is the part that has genuinely been automated. The part that was never the artefact — the judgement, the political work of alignment, the accountability — is untouched, and if anything it now occupies a larger share of what the role actually consists of, because the busywork that used to camouflage its absence has been stripped away.

The steelman for the opposite view deserves a hearing: perhaps, as the tools improve, models will get better at reading the weak signals too, and the judgement gap will narrow. It may narrow. But the accountability gap will not, because it is not a capability gap at all — it is a question of who answers for the outcome, and that will remain a human seat for as long as the consequences land on humans. The programme manager whose value was formatting should indeed be worried. The programme manager whose value was judgement has, this year, simply had the formatting taken off their plate. What they do with the returned hours is the whole question — and the ones spending it in the room rather than on the document are the ones who will still be needed when the novelty wears off.