


BEYOND PROMPT FLUENCY

The AI Skill Shift Is From Producing Information to Designing Decisions

What programme professionals must learn when fluent output becomes cheap and trustworthy action remains hard

Giovanni Leonardi

November 2024



The AI Skill Shift Is From Producing Information to Designing Decisions

Programme · No. 265

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

“The premium is moving from producing information to designing the conditions under which information can be trusted and acted upon.”

1. The report that took minutes and saved nothing

A programme office recently tested a generative AI assistant on a familiar piece of work: turning 18 workstream updates into the weekly executive report.

The old process absorbed most of two days. Analysts chased missing inputs, reconciled inconsistent wording, copied milestones into slides and debated whether one risk belonged under technology or operations. The assistant produced a polished first draft in 11 minutes.

For a moment, the result looked like a breakthrough.

Then the programme director asked three questions. Which dates had changed since the previous week? Which statements were supported by evidence rather than confidence? Which dependency would stop the critical path first?

The draft could not answer reliably because the underlying updates could not answer reliably. It had made weak information fluent. The team saved time on composition and none on judgement.

That episode captures the skill shift facing programme professionals in 2024. The important change is not that we must all learn to write better prompts. It is that the traditional value of collecting, formatting and transmitting information is collapsing faster than our operating models are changing.

AI can produce the artefacts of programme management at remarkable speed. The professional advantage now lies in designing the evidence, decisions and controls that make those artefacts useful.

2. Information coordination is no longer a defensible identity

Programme work has always contained a large amount of information coordination. Plans, status reports, risks, actions, meeting packs, minutes and dependency logs create the shared picture through which leaders govern complex delivery.

That work matters. But much of its labour has been mistaken for its value.

The value is not the report. It is the earlier recognition that a decision is needed.

The value is not the meeting note. It is a durable record of what was decided, by whom, on what evidence and with which consequence.

The value is not the integrated plan. It is the exposure of assumptions and dependencies before they become delays.

Generative AI makes this distinction impossible to ignore. Summarising, reformatting, drafting, classifying and comparing are becoming cheap. A professional role built mainly around those activities will feel increasingly busy and increasingly difficult to justify.

The premium is moving from producing information to designing the conditions under which information can be trusted and acted upon.

This is not a prediction about some distant autonomous future. The shift is already visible in the tools available to teams now. The organisations moving fastest are not removing programme professionals wholesale. They are asking why experienced people still spend Friday assembling material that software can draft, while Monday's governance meeting still avoids the decision that matters.

3. Prompt fluency is useful and radically insufficient

The most common capability response has been to train people to use generative AI tools: how to write prompts, request a format, iterate a draft and check an output.

That is a sensible beginning. It is not a professional development strategy.

Prompt fluency improves the interaction between a person and a model. Programme performance depends on a much larger system: source quality, access, confidentiality, process design, decision rights, assurance, adoption and accountability. A brilliant prompt cannot compensate for undefined milestones, contradictory baselines or a risk process that rewards optimism.

The more durable skills sit around the tool.

3.1 Frame the question before seeking the answer

AI systems are very good at responding to the task they are given, including the wrong task.

Programme professionals need to become more exact about problem framing. What decision will this analysis support? Which constraints are fixed? What evidence would change the conclusion? What time horizon matters? What is outside scope? Who bears the consequence if the answer is wrong?

In the reporting example, “summarise these updates” was easy and of limited value. “Identify changes against the approved baseline, show the evidence for each change and rank the decisions that cannot wait until next week” was closer to the real need.

The difference is not better wording alone. It requires knowledge of delivery mechanics and governance purpose.

3.2 Design work so that machines and people can both understand it

Many programmes run on tacit conventions. A red rating in one workstream means a missed milestone; in another it means a potential resource issue. “Complete” may mean technically built, accepted by operations or merely handed over for review. Dates are copied across plans without recording their source.

AI exposes this ambiguity because it scales whatever definition it receives.

Professionals must learn to design clearer information structures: explicit definitions, stable identifiers, versioned baselines, evidence links, decision records and consistent ownership. This can sound administrative. In practice, it is what allows automation to operate without inventing coherence.

The skill is not data entry. It is creating a delivery language precise enough for the organisation to share.

3.3 Verify proportionately, not ceremonially

AI-generated work requires verification, but “have a human check it” is not an adequate control. If the human repeats the whole task, the efficiency disappears. If the human scans for plausibility, confident errors can pass.

Verification must be designed according to consequence.

- Low-consequence drafting may need a quick review for tone and obvious error.

- Analysis that informs a funding or sequencing decision should be checked against named source evidence.
- A recommendation affecting safety, regulation, employment or contractual position requires appropriate specialist judgement.
- Repeated automated workflows need sample testing, exception monitoring and clear failure handling.

The professional skill is to decide what must be proven, which evidence is sufficient and where human authority cannot be delegated.

3.4 Reconfigure the workflow, not merely accelerate a step

The weakest AI adoption inserts a tool into an existing process and celebrates the time saved in one activity. The surrounding queues remain.

A status draft that takes 11 minutes still creates little benefit if inputs arrive late, evidence is inconsistent and the decision forum meets only once a month. A risk summary generated instantly does not help if risk owners lack authority to act. Faster minutes do not improve governance if decisions remain ambiguous.

Programme professionals need workflow-design skills: map where information originates, where it waits, how it changes, which decision consumes it and what action follows. AI should remove or reshape whole hand-offs, not simply make one document cheaper.

3.5 Lead adoption as a change in authority

AI adoption is not primarily a technology rollout. It changes who prepares information, who interprets it, whose expertise is visible and who is accountable for error.

An analyst asked to automate a report may hear that their judgement is not valued. A senior manager may distrust an output they did not watch someone produce. A specialist may fear that a general model is being allowed to cross a professional boundary. A team may quietly keep the old manual process while also performing the new one.

Programme professionals are well placed to lead this transition, but only if they treat it as an operating-model change. Roles, controls, incentives, learning and escalation all need redesign. Training people on a tool without changing those conditions creates demonstration, not adoption.

4. What the textbooks got right—and what they could not anticipate

There is a strong counterargument that the fundamentals of programme management have not changed. We still need clear outcomes, accountable leadership, integrated planning, risk management, benefits, stakeholder engagement and disciplined governance. Technology fashions come and go; judgement and relationships endure.

That argument is largely correct.

The mistake is to conclude that enduring fundamentals mean enduring professional practices. The purpose of navigation is stable; the instruments can still transform the navigator's work.

A programme professional who understands dependencies, incentives and decision dynamics is more valuable with AI, not less. But that value appears only if the person moves up the chain of reasoning. If they continue to spend most of their time producing summaries, maintaining slideware and policing templates, the enduring fundamentals will not protect the role.

The textbooks prepared us to manage complexity through plans, governance and communication. They did not prepare us for a world in which credible-looking communication can be generated instantly, analysis can be repeated at negligible cost and the boundary between assistance and delegated action is beginning to blur.

The response is not to discard the discipline. It is to practise it at a deeper level.

5. A more honest capability model

The current rush to build “AI skills” often produces a list of courses rather than a view of professional capability.

A stronger model distinguishes four levels.

Level	Capability	Evidence in practice
Tool use	Draft, summarise, compare and structure work safely	Time saved without unreviewed disclosure or obvious quality loss
Work design	Define inputs, rules, hand-offs and exceptions	A whole process becomes shorter, clearer or more reliable
Decision design	Connect evidence to authority, options and consequence	Governance produces earlier, traceable decisions

Level	Capability	Evidence in practice
System stewardship	Monitor quality, behaviour, adoption and control over time	Automation remains useful as data, people and conditions change

Most organisations are concentrating on the first level. The competitive and professional difference will sit in the next three.

Consider a composite transformation with 14 workstreams and 220 active risks. The programme office spends 65 hours each week collecting updates, normalising language and preparing packs. An initial AI pilot reduces document preparation by 24 hours but does not change the governance cycle.

A second redesign makes each risk carry a clear cause, event, consequence, evidence link, review date and decision threshold. Automated analysis now identifies duplicate risks, missing evidence and clusters of dependency. The programme office stops producing three separate risk views and prepares one decision queue. The weekly risk meeting falls from two hours to 45 minutes, and six cross-workstream interventions are made before their trigger dates.

The benefit did not come from superior prose. It came from better risk design, better data and a governance forum organised around action.

That is the kind of evidence a capability programme should seek. Course completion and prompt libraries are weak proxies.

6. What professionals should learn now

The learning agenda should combine practical tool use with disciplines that many programme roles have allowed to atrophy.

- Problem framing: translating a broad concern into a decision, constraint and testable question.
- Information architecture: designing definitions, provenance, identifiers, baselines and evidence relationships.
- Critical evaluation: testing an output for source, assumption, omission, contradiction and consequence.
- Workflow design: removing queues and hand-offs around the automated step.
- Control design: setting permissions, review levels, monitoring, escalation and accountability.

- Commercial judgement: understanding provider dependency, data terms, intellectual property and exit.
- Change leadership: reshaping roles and routines while retaining trust and professional responsibility.
- Facilitation under abundance: helping leaders decide when analysis and options can be generated faster than attention can absorb them.

These skills cannot be developed through classroom instruction alone. Teams need supervised use on real work, with baseline measures and retrospective review.

A practical development cycle can be simple:

1. Choose a recurring programme task with a clear consumer and measurable effort.
2. Define the decision or outcome the task is meant to support.
3. Improve the underlying data and rules before automating the output.
4. Use AI within an agreed confidentiality and review boundary.
5. Measure cycle time, error, rework and decision quality.
6. Record what the human added that the tool could not.
7. Redesign the role around that higher-value contribution.

The final step matters most. If saved time is immediately filled with more reporting, the organisation has automated labour without developing capability.

7. The profession's uncomfortable choice

Programme professionals have often been the people who make complexity legible. That remains a worthy role. But legibility can no longer mean producing more polished descriptions of the work.

It must mean making the system easier to decide within.

This requires a change in professional confidence. Our authority cannot rest on ownership of the template, familiarity with the reporting cycle or the ability to assemble the executive pack under pressure. It must rest on seeing how outcomes, evidence, dependencies, incentives and decisions connect—and on redesigning those connections when they fail.

AI will make some activities dramatically faster and some roles harder to defend. It will also create more need for people who can frame ambiguous problems, distinguish evidence from fluency, design controls and lead changes in how work gets done.

The honest account is therefore neither reassuring nor catastrophic. The core discipline survives, but the comfortable middle of information coordination does not.

The programme professional of value in the AI era will not be the person who can produce the answer fastest. It will be the person who knows which answer can be trusted, which decision it should trigger and how the surrounding system must change for that decision to matter.