

BEYOND THROUGHPUT

When AI Generates Faster Than You Can Evaluate

A Practitioner's Perspective on the Evaluation Bottleneck in AI-Accelerated
Portfolio Delivery

Giovanni Leonardi

September 2025



When AI Generates Faster Than You Can Evaluate

Portfolio · No. 116
September 2025

“The scarce resource in the AI-accelerated portfolio is not production capacity but evaluative attention, and no organisation has yet learned how to budget for it.”

1. The New Bottleneck

For as long as portfolio management has existed as a discipline, the constraint has been production. There were never enough people, never enough time, never enough capacity to pursue every worthwhile initiative. Portfolio governance existed, in large part, to ration that scarce capacity — to ensure the organisation's limited delivery resources were directed at the highest-value work.

That constraint is shifting. Not disappearing — human delivery capacity remains finite — but shifting, because AI tools are now capable of producing certain categories of work at a speed and volume that the portfolio governance system was never designed to absorb. Business cases that once took weeks can be drafted in hours. Options analyses that required dedicated analyst time can be generated on demand. Technical feasibility assessments, market scans, risk registers, benefits models — the raw materials of portfolio decision-making — are becoming available at a pace that outstrips the organisation's ability to evaluate them.

The pattern I have observed is this: the portfolio office receives more material, of reasonable quality, than it can meaningfully assess. The response is almost always one of two extremes. Either the evaluation process slows to a crawl — committees defer decisions because they cannot review everything — or it accelerates dangerously, with approvals granted on the basis of a quick scan rather than genuine scrutiny. Both are failure modes. The first wastes the speed advantage entirely. The second converts it into risk.

2. Why This Is Not a Process Problem

The instinctive response from portfolio leaders is to treat this as a process problem. Speed up the evaluation cycle. Add more reviewers. Streamline the approval workflow. These are reasonable impulses, and they are all insufficient, because they misdiagnose the constraint.

The constraint is not the speed of the evaluation process. It is the depth of evaluative attention that a human decision-maker can bring to bear. A portfolio board that reviews ten business cases in a morning is not doing the same work as one that reviews three. The quality of scrutiny degrades with volume — not because the board members are lazy or careless, but because genuine evaluation of a strategic investment requires the kind of slow, contextual thinking that cannot be parallelised or accelerated without loss.

This is the asymmetry that AI-accelerated portfolios must confront: the production of portfolio artefacts can be scaled with AI, but the evaluation of those artefacts remains stubbornly human and stubbornly slow. And evaluation is where the real value of portfolio management lives. An organisation that can generate fifty business cases but can only genuinely evaluate ten has not gained a fivefold advantage. It has gained a screening problem.

The scarce resource in the AI-accelerated portfolio is not production capacity but evaluative attention, and no organisation has yet learned how to budget for it.

3. Three Patterns Worth Watching

In my experience, three patterns are emerging as organisations begin to grapple with this mismatch. None is fully proven. All are instructive.

The tiered evaluation model. Some portfolio offices are experimenting with a multi-tier approach: AI-generated artefacts pass through an initial automated quality and consistency check, then a rapid human screen (is this worth serious evaluation?), and only then enter the full governance process. The logic is sound — it uses AI to filter AI output before it reaches human evaluators — but the risk is obvious. If the automated screen is too generous, the bottleneck simply moves one stage downstream. If it is too aggressive, genuinely valuable proposals are filtered out before a human ever sees them. Calibrating this screen is itself a human-judgement task, and it requires portfolio professionals who understand both the AI's tendencies and the organisation's strategic context.

The evaluation budget. A more radical approach treats evaluative attention as a finite resource and budgets it explicitly. The portfolio office declares: we will evaluate N proposals per quarter, regardless of how many are generated. This forces prioritisation upstream — sponsors must compete for evaluation slots, which concentrates organisational attention on the proposals that have the strongest pre-screening case. The discomfort this creates is significant. Leaders accustomed to submitting everything for consideration resist being told that their proposal will not even be evaluated until next quarter. But the discomfort is honest — it makes explicit a constraint that was always present but previously hidden behind the slower pace of production.

The evaluator-in-the-loop model. Rather than separating production and evaluation into sequential phases, some organisations are embedding evaluators into the production process itself. The business case is not generated in isolation and then submitted; it is developed iteratively with a designated evaluator who shapes the analysis, challenges assumptions, and builds judgement as the artefact takes shape. By the time the case

reaches the portfolio board, the evaluator can vouch for its rigour — not because they rubber-stamped it, but because they participated in its construction. This approach is resource-intensive, but it converts the evaluation problem from a bottleneck into a collaboration.

4. What This Means for Portfolio Leaders

The implications for portfolio management practitioners are significant and uncomfortable.

First, the value proposition of the portfolio office shifts. If AI can generate the analytical artefacts that portfolio offices have traditionally produced or commissioned, the office's value no longer lies in orchestrating that production. It lies in the quality of its evaluative judgement — its ability to distinguish a compelling case from a plausible one, to see the risks that a well-constructed business case obscures, and to maintain strategic coherence across a portfolio of initiatives that is growing faster than the organisation's capacity to deliver. This is a more demanding role, not a less demanding one, and it requires a different skill profile: less analytical production, more strategic judgement.

Second, governance cadences need rethinking. Quarterly portfolio reviews made sense when the input was three months' worth of slowly-assembled proposals. They make less sense when the same volume of proposals can be generated in a week. The cadence must either accelerate — with all the risks of shallower evaluation — or the intake process must decouple from the governance cycle, allowing proposals to enter a managed queue rather than waiting for the next scheduled review.

Third, the relationship between sponsors and the portfolio office changes. When producing a business case was expensive in time and effort, sponsors self-selected — only initiatives with genuine backing were worth the investment of assembling a case. When AI makes case production cheap, the self-selection mechanism weakens. More proposals arrive, of more variable strategic merit, and the portfolio office must develop new mechanisms for managing that volume without becoming a gatekeeper that stifles innovation or a bottleneck that frustrates sponsors.

5. The Deeper Question

Beneath all of this lies a question that portfolio management has never had to ask quite so directly: what is evaluation actually for?

If evaluation is a quality check — does this business case meet our standards? — then it can, in principle, be partially automated. AI can check for consistency, completeness, and alignment with stated criteria. But if evaluation is a judgement — is this the right thing to do, given everything we know about our strategy, our capacity, our risk appetite, and the hundred other things competing for the same resources? — then it cannot be automated, because it requires exactly the kind of contextual, political, and experiential reasoning that makes human judgement irreplaceable.

The organisations that will manage AI-accelerated portfolios most effectively are those that are honest about which parts of their evaluation process are quality checks and which are judgements — and that protect the space for judgement even as they automate the checks. The temptation to collapse the distinction — to treat everything as a check, because checks can be scaled — will be strong. It should be resisted.

Portfolio management has always been, at its core, an exercise in attention allocation: directing the organisation's limited capacity toward its highest-value opportunities. The AI era does not change that purpose. It makes it harder, because the volume of plausible opportunities has increased while the capacity for genuine evaluation has not. The discipline's future depends on whether it can adapt its methods to that new reality without losing the judgement that gives those methods their meaning.