

BEYOND THE PIPELINE

Portfolio Management When AI Can Generate Faster Than Humans Can Evaluate

A Practitioner's Perspective on the Evaluation Bottleneck and the Future of
Portfolio Governance

Giovanni Leonardi

November 2025



Contents

1. Executive Summary	3
2. The Old Bottleneck	3
3. The New Imbalance	4
4. The Five Responses	4
5. The Structural Forces	5
6. What This Reveals	6
7. The Path Forward	7

Portfolio Management When AI Can Generate Faster Than Humans Can Evaluate

Portfolio · No. 117
November 2025

“The portfolio office that measures its health by the number of initiatives in flight has already lost — the scarce resource is no longer ideas or even execution capacity, but the organisational attention required to judge whether any of it is worth doing.”

1. Executive Summary

Across sectors, a new pattern is emerging in portfolio management: the rate at which AI systems can generate proposals, business cases, technical designs, and implementation plans has begun to outstrip the organisation's capacity to evaluate them. Portfolio offices that once struggled with a pipeline of twenty or thirty initiative proposals per quarter are now processing hundreds, many of them AI-assisted or AI-originated, each arriving with a level of polish and analytical rigour that makes cursory dismissal difficult. The result is not a crisis of quality but a crisis of attention — and it is exposing structural weaknesses in how organisations make portfolio-level decisions that long predate the current wave of AI capability.

This essay examines why this imbalance persists, what forces sustain it, and what it reveals about the gap between how organisations intend to manage their portfolios and how portfolio decisions are actually made.

2. The Old Bottleneck

For most of the history of portfolio management as a discipline, the binding constraint has been generation. Coming up with viable initiative proposals — well-researched, properly costed, strategically aligned, supported by a credible business case — required significant human effort. A single business case might take weeks to develop. Strategic options analysis consumed months. The portfolio pipeline was naturally rate-limited by the intellectual and administrative labour of producing proposals that met the bar for serious consideration.

This natural rate-limiting had a secondary benefit that was rarely acknowledged: it gave the evaluation function time. A portfolio board reviewing twelve proposals per quarter could give each one genuine scrutiny. Board members could read the materials,

interrogate the assumptions, compare the initiative against strategic priorities, and form a considered judgement. The scarcity of proposals created the conditions for thoughtful evaluation.

The portfolio management frameworks that emerged during this period reflect this reality. They assume a manageable volume of proposals, each requiring detailed human review. Stage-gate models, weighted scoring matrices, strategic alignment assessments – all presuppose that the evaluation function can keep pace with the generation function. The entire discipline was built around the implicit assumption that producing good proposals was harder than judging them.

That assumption no longer holds.

3. The New Imbalance

What has changed is not that AI has replaced human judgement in proposal generation – it has not, or at least not yet. What has changed is that AI has dramatically reduced the cost and time required to produce a proposal that looks evaluable. An AI-assisted team can generate a business case in hours that would previously have taken weeks. It can produce strategic options analyses, financial models, risk assessments, and implementation roadmaps at a pace that renders the old pipeline assumptions obsolete.

The quality question is genuinely complex. Many AI-assisted proposals are substantively strong – they draw on broader data, consider more variables, and present more rigorous quantitative analysis than their purely human-authored predecessors. But some are sophisticated surfaces over shallow foundations: analytically polished documents whose underlying assumptions have not been subjected to the kind of adversarial scrutiny that a slower, more laborious authoring process would have imposed. The problem for the portfolio board is that distinguishing between these two categories requires exactly the kind of deep, time-intensive evaluation that the volume of proposals now makes impossible.

The numbers tell the story. In conversations with portfolio leaders across financial services, telecommunications, and government, a consistent pattern emerges: the volume of initiative proposals entering portfolio governance has increased by a factor of three to five since AI-assisted business case development became widespread. The evaluation capacity of portfolio boards has not changed. The same twelve people meet for the same three hours on the same monthly cycle, now facing five times the volume.

The portfolio office that measures its health by the number of initiatives in flight has already lost – the scarce resource is no longer ideas or even execution capacity, but the organisational attention required to judge whether any of it is worth doing.

4. The Five Responses

Organisations are responding to this imbalance in ways that fall into five broad categories, none of which is fully satisfactory.

Raising the bar. Some portfolio offices have responded by increasing the requirements for proposal submission — more detailed business cases, more rigorous financial modelling, more extensive stakeholder consultation before a proposal reaches the board. The logic is that higher entry requirements will reduce volume to manageable levels. In practice, AI-assisted teams clear higher bars as easily as lower ones, and the main effect is to disadvantage teams without AI capability while doing little to reduce the queue.

Delegation and tiering. Others have introduced tiered evaluation, with only the largest or most strategic proposals reaching the full portfolio board and smaller initiatives delegated to departmental governance. This is sensible in principle but introduces a new risk: the cumulative resource and attention cost of many small initiatives can exceed that of a single large one, and delegated governance often lacks the cross-portfolio visibility to spot conflicts, duplications, or competing demands on shared capabilities.

AI-assisted evaluation. A growing number of organisations are using AI to evaluate the AI-generated proposals — automated scoring against strategic criteria, algorithmic comparison of financial projections, pattern-matching against historical initiative outcomes. This accelerates the filtering process but raises an uncomfortable question about what portfolio governance is actually for. If both the generation and evaluation of proposals are AI-driven, the human role reduces to ratifying the output of a system whose internal logic is opaque to most board members.

Moratorium and batching. Some organisations have imposed intake windows — proposals accepted only during defined periods, with everything outside the window deferred. This restores the time pressure that the old generation bottleneck provided naturally, but it sits uneasily with the stated ambition of agile, responsive portfolio management. It is, in effect, an admission that the organisation cannot govern continuous flow and must retreat to batch processing.

Ignoring the problem. The most common response, if the least discussed, is to continue operating as though the volume has not changed. Portfolio boards process proposals faster, which in practice means processing them less thoroughly. Review sessions become approval sessions. The governance ritual continues, but the scrutiny that justified it has quietly evaporated.

5. The Structural Forces

The persistence of this pattern — across sectors, across organisational types, across varying levels of AI maturity — suggests that it is not a transitional problem that will resolve as organisations adapt. It is structural, sustained by forces that operate independently of the specific technology.

The generation-evaluation asymmetry is inherent. Generating a plausible proposal is a fundamentally different cognitive task from evaluating one. Generation is convergent — it assembles information into a coherent narrative. Evaluation is divergent — it questions assumptions, imagines failure modes, considers second-order effects, and weighs incommensurable values. AI is exceptionally good at the former and still limited at the

latter, because genuine evaluation requires the kind of contextual judgement, organisational knowledge, and political awareness that remains stubbornly human. Any technology that accelerates generation without equally accelerating evaluation will widen the gap.

Portfolio governance is socially constructed. The portfolio board is not merely an analytical function — it is a political arena in which competing interests negotiate resource allocation. The value of portfolio governance lies partly in the conversation it forces: the debate about priorities, the surfacing of conflicts, the negotiation of trade-offs. This conversation cannot be accelerated without being degraded, because it depends on human relationships, institutional memory, and the kind of trust that develops only through repeated interaction. Faster throughput through the governance process is not the same as faster governance.

The incentive structure rewards generation. In most organisations, the people who generate proposals are rewarded for volume and ambition. Business unit leaders who bring forward more initiatives are seen as dynamic and strategic. The people who evaluate proposals bear the reputational cost of saying no but receive little credit for maintaining portfolio discipline. This asymmetry has always existed, but AI amplification has made it acute. When generating a proposal is cheap and blocking one is expensive, the portfolio inevitably inflates.

“We have built organisations that reward the production of plans and penalise the exercise of judgement about whether those plans should exist.”

6. What This Reveals

The generation-evaluation imbalance is not, at its core, an AI problem. It is a governance problem that AI has made visible. The organisations struggling with proposal volume in 2025 are, in most cases, organisations that were already struggling with portfolio discipline before AI entered the picture. They approved too many initiatives, spread resources too thinly, failed to kill underperforming programmes, and treated the portfolio as a collection of projects rather than an integrated investment portfolio.

AI has not created these dysfunctions. It has amplified them by removing the natural friction that previously masked them. When it took six weeks to produce a business case, the sheer effort involved imposed a de facto discipline on the pipeline. That discipline was never a governance achievement — it was a side effect of administrative burden. Now that the burden has been lifted, the underlying governance weakness is exposed.

This is, in many ways, the central insight: what we mistook for governance discipline was actually generation friction. The approval chain worked not because it was well-designed but because the pipeline was slow. Speed up the pipeline, and the governance reveals itself as what it always was — a thin layer of review over a process that was never truly

governed.

7. The Path Forward

There is no simple fix for this imbalance, because the imbalance is a symptom of a deeper problem. But the outlines of a more honest portfolio governance model are becoming visible.

First, portfolio boards must shift from evaluating proposals to governing the conditions under which proposals are generated. This means defining strategic corridors — the domains in which new initiatives are welcome — and making clear that proposals outside these corridors will not be considered regardless of their quality. This is uncomfortable because it requires the organisation to say, explicitly and in advance, what it will not do. Most organisations prefer to keep their options open, which is precisely why the portfolio inflates.

Second, the profession must develop evaluation methods that scale without sacrificing the depth that makes evaluation valuable. This does not mean automating evaluation — it means redesigning evaluation to focus human attention on the decisions that genuinely require it: the strategic trade-offs, the resource conflicts, the risk judgements that no algorithm can make well. Everything else — the compliance checks, the financial validation, the alignment scoring — can and should be automated, not to replace governance but to clear the ground so that governance can focus on what matters.

Third, organisations must confront the incentive asymmetry directly. If portfolio discipline matters, then the people who maintain it must be valued as highly as the people who generate proposals. This is a cultural challenge more than a structural one, and it will not be solved by process redesign alone. It requires senior leaders to model the behaviour they want to see: to celebrate the decision not to proceed as much as the decision to invest, and to treat portfolio restraint as a strategic capability rather than a failure of ambition.

The generation-evaluation imbalance will not resolve itself. AI will continue to lower the cost of producing plausible proposals. The organisations that thrive will be those that recognise the bottleneck has moved — from generation to evaluation, from production to judgement, from ambition to discipline — and redesign their governance accordingly.