

BEYOND FALSE PRECISION

The AI Business Case Problem — Pricing Uncertainty When the Technology Moves Quarterly

Investment Governance for a Technology That Won't Hold Still

Giovanni Leonardi

May 2025



The AI Business Case Problem — Pricing Uncertainty When the Technology Moves Quarterly

Portfolio · No. 115
May 2025

“A process that feels comfortable when funding AI in the middle of 2025 is almost certainly providing false comfort — packaging ignorance as confidence, pretending to know what cannot yet be known.”

1. The Familiar Ritual, Broken

A transformation director presents a twelve-month AI business case to the investment committee. The model is impeccable: development costs benchmarked against two vendors, licensing fees confirmed, projected efficiency gains triangulated through three departments. It took six weeks to build and a fortnight of steering-group review. By the time it reaches the approval meeting, the foundation model it was costed against has been superseded, the inference pricing has dropped by forty per cent, and a competitor has released a capability that achieves the same outcome with a fraction of the custom development. The committee approves. The number is already wrong.

This is not a failure of financial modelling. It is the collision between a governance instrument built for stable investments and a technology that rewrites its own economics every quarter. We have spent decades perfecting the business case on the assumption that the solution holds still long enough to cost it. Artificial intelligence breaks that assumption so completely that the traditional case has become a form of institutional theatre — a document that provides the appearance of rigour without the substance of truth.

The tension between short-horizon delivery and long-horizon investment planning is familiar to anyone who has managed a transformation portfolio. AI does not create this tension. It accelerates it beyond the point where the traditional instruments can cope.

2. Two Breakdowns, One Root Cause

The business case rests on a foundational premise: that you can specify what you are building, estimate what it will cost, and project what it will return. AI undermines every element, but the damage flows from two distinct mechanisms that compound one another.

The first is capability acceleration. Foundation models are not improving on the timescale of annual planning — they are improving on the timescale of quarterly product releases. What required a bespoke fine-tuned model six months ago can be accomplished with a general-purpose API call today. The custom natural-language processing pipeline that justified a team of eight in January may be replicable with a prompt and a retrieval layer by July. This pattern has repeated across text generation, code assistance, image analysis, and document processing throughout 2024 and into 2025. The implication for the business case is brutal: today's build is tomorrow's commodity. A case that assumes a fixed capability at a fixed cost is pricing a photograph of something that is moving.

The second is price collapse. Inference costs have fallen not by the incremental percentages familiar from cloud computing but by orders of magnitude. A million tokens that cost sixty dollars eighteen months ago may cost two dollars today, and the trajectory shows no sign of stabilising. When the API call that underpins your projected return becomes ten times cheaper between approval and first deployment, the entire financial model shifts — sometimes favourably, sometimes by rendering the custom investment unnecessary altogether.

These forces interact. Capability improvements reduce the need for custom development. Price collapse makes the remaining development cheaper but also makes the wait-and-buy alternative progressively more attractive. The result is a technology where the optimal investment decision changes faster than the governance cycle can process it.

The problem is not that the business case lacks rigour. It is that rigour applied to fictional inputs produces a precisely wrong answer.

3. The Governance Reflex That Makes It Worse

The instinctive organisational response to uncertainty is to demand more certainty from the business case. More detailed cost modelling. More rigorous benefit quantification. Longer approval chains. This reflex is precisely backwards. A beautifully modelled case built on assumptions that will be invalid in ninety days is not a rigorous case. It is an expensive fiction.

I have watched this pattern across a dozen organisations now entering their second or third year of enterprise AI investment. The governance function, understandably anxious about the sums involved, tightens the business case requirements. Programme teams, knowing the numbers are speculative, invest weeks in producing forecasts they privately regard as unreliable. The investment committee, reassured by the apparent precision, approves. Six months later the assumptions have shifted so far that the programme is either dramatically over-delivering against a case that was too conservative, or pursuing an approach that a later model release has rendered unnecessary. In neither outcome did the business case serve its intended function.

The strongest objection to what follows is that organisations cannot fund AI initiatives without financial discipline — that the business case, however imperfect, imposes

necessary accountability. This objection is correct, and nothing in the alternative I propose abandons accountability. The question is not whether to impose discipline but what form discipline should take when the underlying economics move quarterly. A business case designed for a factory expansion — where the machinery, the costs, and the market are knowable within useful bounds — is a different instrument from what AI investment requires. Deploying the wrong instrument is not discipline. It is comfort.

4. Funding What You Can See

The adaptation draws on a principle that venture capital has applied for decades: fund in increments matched to the pace at which uncertainty resolves. What is new is bringing that principle inside the enterprise, where governance cultures have been optimised for predictability rather than learning.

In practice this means structuring AI investments as funding stages of eight to twelve weeks, each with a defined learning objective rather than a fixed deliverable. The first tranche funds a proof of concept that tests whether a foundation model can perform a specific task at the required accuracy. The gate at the end asks not "has the team delivered a product?" but "have we learned enough to justify the next stage?" The answer might be proceed. It might be pivot — the proof of concept revealed that a different application of the same capability is more valuable. It might be kill — the technology cannot meet the need at this cost, or a commercial product has emerged that makes the custom build redundant.

Three mechanics make this operational:

Kill and pivot criteria, defined before the tranche begins. Each funding stage specifies in advance the conditions under which the initiative stops or redirects. This is where the governance discomfort is sharpest. Traditional business cases rarely contain their own termination conditions — the assumption is that the approved case will be delivered. Venture-style tranching inverts this: the default is that an initiative ends unless it earns continuation.

Build-versus-wait as a priced option. For any AI capability there is a live question: build now, or wait for the technology to mature and the price to drop? This is a real options problem, and it can be priced — imperfectly but usefully. The option value of waiting is a function of how fast capability is improving in the relevant domain, how fast prices are falling, and how much competitive value accrues to early deployment. A customer-facing AI capability in a contested market may justify the premium of building now, because the learning and market position gained have independent worth. An internal process automation with no competitive clock ticking may be better served by waiting six months for a superior, cheaper solution. This is not a binary decision but a continuously repriced option, revisited at every tranche gate.

Portfolio-level probability rather than initiative-level certainty. No single AI investment funded in tranches will carry the clean ROI narrative that an investment committee expects. The appropriate unit of governance is the portfolio: across ten initiatives, each funded in short stages with explicit kill criteria, the expectation is that the portfolio returns value even though individual bets will be killed, pivoted, or superseded. This is the

fundamental shift — from demanding certainty at the initiative level to managing probability at the portfolio level.

5. The Discomfort Is the Point

Enterprise governance has been optimised for a world in which senior leaders approve large investments on the strength of detailed, confident forecasts. Venture-style AI funding asks them to approve smaller investments on the strength of structured uncertainty — to fund the next stage of learning without knowing where that learning will lead, and to treat the willingness to stop as a feature of the model rather than a failure of commitment.

This discomfort is not a defect. It is evidence that the governance mechanism is doing honest work. A process that feels comfortable when funding AI in the middle of 2025 is almost certainly providing false comfort — packaging ignorance as confidence, pretending to know what cannot yet be known. The honest position is that we are investing in a technology whose economics, capabilities, and competitive implications are shifting faster than any annual cycle can track, and that the appropriate response is governance whose cadence matches the technology's pace of change.

The business case is not dead. But the business case as a single, upfront, multi-year commitment against stable assumptions is fiction when the technology moves quarterly. The organisations that will navigate this well are those that replace the comfort of false precision with the discipline of iterative, evidence-based funding — and discover in that discipline a governance instrument better suited to the technology they are trying to deploy.