

BEYOND THE STAGE GATE

Portfolio Investment Models for AI — When the ROI Curve Is Exponential

A Practitioner's Perspective on Why Traditional Stage-Gate Governance Kills AI Investments Before They Mature

Giovanni Leonardi

November 2024



Portfolio Investment Models for AI — When the ROI Curve Is Exponential

Portfolio · No. 125

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

“An AI investment that looks like a failure at month six and a marginal performer at month twelve can become the highest-returning asset in the portfolio by month twenty-four — but only if it survives the governance process designed to kill exactly this pattern.”

1. The Governance Trap

Organisations across every sector are now committing significant capital to artificial intelligence. The investment cases are compelling: operational efficiency through automation, improved decision-making through predictive analytics, new revenue streams through intelligent products and services. The strategic logic is sound. And yet a troubling pattern is emerging in how these investments are governed, measured, and — too often — terminated.

The problem is not that organisations are investing in the wrong things. It is that they are applying the wrong investment model to the right things. Traditional portfolio governance assumes a roughly linear relationship between investment and return. Capital is deployed in stages, value is expected to accrue incrementally, and each stage gate asks the same question: is this investment delivering returns commensurate with its cost? For most investments — infrastructure upgrades, process improvements, system replacements — this model works well enough. For AI investments, it is structurally destructive.

AI investments do not follow linear return curves. They follow exponential ones. The early stages consume capital while producing modest, sometimes negligible, visible returns. The data must be gathered, cleaned, and structured. The models must be trained, tested, and refined. The organisation must develop the capability to interpret and act on AI-generated insights. This foundational work is expensive, unglamorous, and produces almost nothing that a traditional stage-gate review would recognise as value.

The consequence is predictable: AI investments are killed at precisely the point where they are about to become valuable.

2. Why Linear Models Fail

To understand why the mismatch is so damaging, it is worth examining what a traditional portfolio investment model actually measures and how it makes decisions.

A standard stage-gate process evaluates investments at predetermined intervals against a set of criteria: financial returns to date, projected returns against the original business case, risk profile, strategic alignment, and resource consumption. Investments that fall below threshold on these criteria are escalated, restructured, or terminated. The model is designed to be disciplined — to prevent good money chasing bad, to enforce accountability, and to ensure that the portfolio as a whole delivers against its aggregate targets.

These are sensible objectives. The problem is that the measurement framework embeds an assumption about the shape of value creation that does not hold for AI investments. The framework expects to see value accumulating steadily from early in the investment lifecycle. It expects that a reasonable proportion of the ultimate return should be visible by the time the investment reaches its first or second stage gate. And it expects that the trajectory of returns at any given point is a reliable predictor of future performance.

None of these assumptions hold for AI. The value creation pattern looks more like this:

Investment Phase	Traditional Investment	AI Investment
Months 1–6	10–15% of projected value visible	Near-zero visible value; foundational data and model work
Months 6–12	30–40% of projected value visible	5–10% of projected value; early model outputs, limited adoption
Months 12–18	50–60% of projected value visible	15–25% of projected value; model refinement, growing accuracy
Months 18–24	70–80% of projected value visible	50–80% of projected value; exponential acceleration as adoption scales
Months 24–36	90–100% of projected value	100–300% of projected value; compounding network effects and capability gains

The traditional investment is comfortably inside the governance framework's tolerance at every review point. The AI investment is outside it at every review point until the moment it overtakes everything else in the portfolio. A governance process that cannot distinguish between these two patterns will systematically eliminate the second.

3. The Compounding Effect

The exponential shape of AI returns is not arbitrary. It is driven by specific dynamics that are inherent to how AI systems create value.

Data flywheel effects. AI models improve as they process more data. The more transactions a fraud detection model analyses, the more accurate it becomes. The more customer interactions a recommendation engine processes, the more precisely it targets. This creates a compounding loop: better performance drives greater adoption, which generates more data, which drives further improvement. The value curve steepens with use.

Capability transfer. Once an organisation builds foundational AI capabilities — data infrastructure, model development skills, deployment pipelines, change management practices — those capabilities can be applied to new use cases at a fraction of the original cost. The first AI initiative is expensive. The second is significantly cheaper. The fifth is transformatively cheap. This portfolio-level dynamic is invisible to a governance framework that evaluates each initiative in isolation.

Adoption thresholds. AI-generated insights are only valuable if the organisation acts on them. In the early stages, users are sceptical, processes are not adapted, and the AI outputs sit alongside human judgement rather than informing it. There is typically a threshold — often around the twelve-to-eighteen-month mark — where confidence in the system reaches a tipping point and adoption accelerates rapidly. Before the threshold, returns are minimal. After it, they compound.

1. These three dynamics interact. Data flywheel effects make the models better, which drives adoption past the threshold, which generates more data, which further improves the models, which enables capability transfer to new domains.
2. The result is not a linear progression but an s-curve — flat at the bottom, steep in the middle, and eventually flattening at a much higher level than linear models would predict.
3. The governance implication is profound: the investment's performance at any point on the flat part of the curve tells you almost nothing about its ultimate value.

4. What a Fit-for-Purpose Model Looks Like

If traditional stage-gate governance is structurally wrong for AI investments, what should replace it? The answer is not to abandon governance — that would be irresponsible. It is to redesign the governance framework so that it can accommodate non-linear value creation while still maintaining investment discipline.

The model I would advocate rests on four principles:

Milestone-based rather than time-based gates. Instead of reviewing AI investments on a fixed calendar — quarterly, semi-annually — review them against capability milestones that reflect the actual dynamics of AI value creation. Has the data infrastructure reached the quality threshold the model requires? Has model accuracy crossed the adoption tipping point? Are capability transfer opportunities identified and costed? These questions are more diagnostic than "what is the ROI at month six?"

Leading indicators rather than lagging returns. Traditional stage gates focus on financial returns — a lagging indicator. For AI investments in their early phases, the right metrics are leading indicators of future return: data quality and volume trends, model accuracy

trajectories, user adoption rates, time-to-insight improvements, and the pipeline of potential capability transfers. An AI investment can be progressing exactly as it should while producing negligible financial return. The governance framework needs metrics that can tell the difference between "not yet delivering" and "never going to deliver."

Portfolio-level evaluation. Because AI capability transfers create portfolio-level value that no individual investment captures, the governance framework should evaluate AI investments partly on their contribution to the organisation's overall AI capability. An initiative that delivers modest direct returns but builds foundational data assets, develops organisational skills, or proves a deployment pattern that accelerates subsequent initiatives is creating real value that an initiative-level review will miss.

Protected incubation periods. For AI investments that pass initial feasibility and alignment criteria, the governance framework should provide a protected incubation period — typically twelve to eighteen months — during which the investment is shielded from standard ROI-based stage gates. During this period, the investment is governed on leading indicators and capability milestones. After the incubation period, the investment transitions to standard portfolio governance, by which point the exponential return dynamics should be visible in the financial data.

The goal is not to exempt AI investments from accountability. It is to hold them accountable for the right things at the right time — rather than measuring them against a value curve they were never designed to follow.

5. The Organisational Prerequisites

Redesigning the governance framework is necessary but not sufficient. Several organisational conditions must be in place for an exponential-aware investment model to work.

The portfolio leadership team must understand the dynamics of AI value creation well enough to defend investments that look like failures under traditional metrics. This is a literacy challenge as much as a governance challenge. If the Chief Financial Officer and the portfolio board do not understand why an AI investment can be performing well while showing poor financial returns, no amount of framework redesign will save it from termination at the first uncomfortable stage gate.

The organisation must be willing to ring-fence AI investment capital from the general portfolio budget. If AI investments compete head-to-head with traditional investments for the same pool of capital using the same evaluation criteria, they will lose every time. The linear investments will always look better at early stage gates. A separate AI investment allocation — sized appropriately, governed distinctly, and measured on its own terms — is the practical mechanism for preventing the governance trap.

Finally, the organisation must accept a different risk profile for AI investments. Traditional portfolio management seeks to minimise variance — to produce a reliable, predictable aggregate return. AI investments are inherently higher-variance: some will fail, some will produce modest returns, and a few will produce extraordinary returns. The

portfolio must be sized and diversified to absorb the failures while capturing the outliers.

6. The Cost of Inaction

Organisations that continue to apply linear governance models to exponential investments will pay a price that compounds over time. They will terminate AI initiatives that were months away from inflection. They will lose the foundational capabilities those initiatives were building. They will watch competitors who tolerated the uncomfortable early phase pull ahead as their AI investments mature and compound.

The most damaging outcome is not any single terminated investment. It is the organisational learning that accumulates: the belief, reinforced by repeated governance decisions, that AI does not deliver. Each killed initiative becomes evidence for the thesis that AI investments are not worth making — a thesis that is self-fulfilling precisely because the governance model ensures that no AI investment survives long enough to prove it wrong.

“The organisations that will capture the full value of artificial intelligence are not the ones making the largest investments. They are the ones whose governance frameworks are sophisticated enough to let those investments mature.”

The investment model is the strategy. Getting it wrong is not a governance failure. It is a competitive one.