



MyConsultancy

GOVERNING AI PORTFOLIOS, NOT PILOTS

From AI experiments to disciplined investment
and scaled value



Manoj Tavarajoo


EXECUTIVE SUMMARY

AI value does not come from running more experiments. It comes from governing the right things, at the right level, with the discipline to concentrate investment and scale what matters.

Most organisations are not doing this. They are approving pilots, funding use cases, and reporting on activity. The result is a portfolio that looks productive and performs poorly. A large number of initiatives. A small amount of scaled value. Significant capital deployed with limited commercial return.

The root cause is structural. Organisations are governing AI at the wrong unit of analysis. A use case solves a local problem. A strategic capability creates repeatable, cross-enterprise value. Governing at use case level produces a catalogue. Governing at capability level produces a portfolio. Very few organisations have made this shift, and the financial consequences are now visible.






“
AI portfolios are capital
allocations to capability.
Not pipelines of use cases.”

The inability to scale AI is, in most cases, a governance failure rather than a technology or talent failure. The organisations generating disproportionate returns are not experimenting more broadly. They are investing with more discipline. They treat AI as a capital allocation decision. They concentrate investment, build reusable foundations, assign clear business ownership for outcomes, and stop what is not working.

This paper sets out the principles, frameworks, and governance mechanisms required to make that shift. It builds on the portfolio discipline and operating model foundations established in *The AI Operating Model Playbook* and the leadership case made in *Leading the AI Transformation*, extending both into the investment governance structures needed to deliver value at enterprise scale.



THE GOVERNANCE PROBLEM

Across industries, organisations are committing substantial capital to AI while struggling to convert that investment into sustained business value. Initiatives are launched, proofs of concept are completed, and models are demonstrated. Few of these efforts progress into scaled, operational capabilities that move the dial on revenue, cost, or competitive position.

The evidence is consistent and pointed. McKinsey's State of AI 2025 found that nearly two-thirds of organisations remain in the experimentation or pilot stage. BCG reports that 60 per cent of companies are generating no material value from their AI investments and 74 per cent have yet to demonstrate tangible returns. MIT's GenAI Divide research found that only 5 per cent of enterprise generative AI pilots achieve meaningful P&L impact.

These are not technology numbers. They are governance numbers. The models work. The data is available. The talent exists. What is missing is the discipline to treat AI investment as a capital allocation to capability rather than a portfolio of individual initiatives each seeking their own approval.

AI is typically managed as a series of projects, a pipeline of use cases, or an innovation programme. Governance reflects this framing. Decisions are made at the level of individual initiatives. Funding is allocated case by case. Ownership sits across business units and technology teams, with no single point of accountability for whether the portfolio, as a whole, is generating value.

The result is that organisations accumulate activity without building capability. Effort is duplicated, standards are inconsistent, and reuse is left to chance. The portfolio appears active. Value does not scale.

More AI activity does not mean more AI value. It often means the opposite.

WHY PILOT CULTURE PERSISTS

Pilot culture is not a behaviour problem. It is a governance design problem. The funding models, incentive structures, and approval processes of most large organisations are built to encourage the start of initiatives, not their completion at scale.

Funding models favour small, contained experiments. Securing approval for a pilot is far easier than committing to a long-term capability. Innovation programmes measure what they reward: the number of ideas tested. Leadership visibility tends to accrue to the executive who launched an initiative, not the one who scaled it into operations and closed it with a benefits case.



Governance mechanisms compound this. In most organisations, AI opportunities enter through multiple channels, are evaluated on local criteria, and are approved by different decision-makers with no shared view of portfolio impact. There is no mechanism for comparing initiatives rigorously across the whole portfolio, no process for concentrating investment deliberately, and no requirement for an approved initiative to demonstrate a credible pathway to scale.

The consequence is a system that optimises for starting. As Deloitte's 2025 AI ROI research found, even executives who acknowledge unclear returns continue to approve new initiatives, driven by the fear of falling behind competitors rather than by a clear view of where the next dollar of AI investment will generate the greatest return.

There is a further signal worth noting. The rise of unsanctioned AI usage across organisations reflects a misalignment between formal governance pathways and operational demand. Employees adopt tools independently when the sanctioned process is too slow or too constrained. This is not only a risk management challenge. It is evidence that the governance model is not fit for the pace and scale of demand it is supposed to manage.

“

Pilot culture persists because the governance model was designed to start things, not to scale them.




WHY IT MATTERS NOW

Two groups of organisations are now emerging, and the distance between them is growing.

The first group treats AI as a portfolio of capital allocations to capability. They make concentrated bets on a small number of high-impact capabilities, embed those capabilities into core operations, and measure outcomes against business performance rather than initiative count. They are building a compounding asset. Each capability funded and scaled lowers the marginal cost of the next. Each platform investment reduces the effort required for every subsequent use case that draws on it.

The second group treats AI as a programme of activity. They spread investment across a wide range of initiatives, many of which will not survive beyond early stages. They report on experiments. They celebrate pilots. They reapprove variations of things they have already tried. Capital is consumed without accumulating. The portfolio grows in length and diminishes in value.

PwC's analysis reinforces the distinction. High-performing organisations apply explicit portfolio logic: allocating investment across short-term efficiency gains, medium-term capability bets, and a small number of transformational commitments, with each phase generating returns that fund the next.



WHERE CURRENT APPROACHES FALL SHORT


Much of the current discussion on AI governance focuses on responsible use, risk classification, and operating model design. Frameworks such as ISO/IEC 42001, the NIST AI Risk Management Framework, and the EU AI Act provide essential structure for managing AI systems responsibly. These are necessary foundations.

What they do not address is how to manage AI investments as a portfolio of strategic capabilities. That is a different discipline, and for most organisations it is a more immediate source of value loss.

Current approaches evaluate initiatives individually rather than collectively. They prioritise on local business unit criteria rather than portfolio impact. They fund projects rather than build persistent capability. They treat reuse as a technical consideration rather than an economic argument. Risk is assessed at the system level, not at the portfolio level.

The result is that organisations cannot answer the questions that matter most at enterprise level. Where is value concentrating in this portfolio? What should not be funded? Which capabilities are being built for reuse, and which are being rebuilt from scratch in each business unit? Who is accountable for outcomes, not just outputs?

Use case approval is not portfolio governance. It is the absence of it. A list of approved use cases is a catalogue. A portfolio has investment logic, concentration decisions, explicit trade-offs, and outcome accountability. The distinction is straightforward, and most organisations have not made it.



McKinsey's research confirms the commercial consequence. Organisations investing at scale in AI governance and capability are significantly more likely to achieve material EBIT impact. Governance investment is not a cost. It is the structural condition for return.

The organisations in the first group are not using different technology. They are making better capital allocation decisions and governing those decisions with discipline.



“

The compounders and the accumulators are separating. The difference is not technology. It is investment discipline.

The scale of the divergence is already measurable. PwC's analysis of senior executives across major economies found that the top fifth of companies are capturing close to three-quarters of the economic value AI is generating. The remaining four-fifths are sharing what is left.

As AI becomes central to how businesses compete and operate, this advantage compounds. Establishing portfolio discipline early creates a structural lead that becomes progressively harder for late movers to close.

CORE PRINCIPLES OF AI PORTFOLIO GOVERNANCE

The shift from use case management to portfolio governance is a shift in discipline. The following principles define what that discipline requires in practice.

Governance is capital allocation

Every AI governance decision is an investment decision. Which initiative deserves capital? Which should be deferred? Which should not be funded at all? A governance function that cannot answer these questions at portfolio level is not governing. It is processing applications.

A portfolio that funds everything cannot scale anything

Organisations that generate disproportionate AI value make concentrated bets on a small number of high-impact capabilities and resource those bets properly. Spreading capital thinly across many initiatives is not prudent diversification. It is the systematic prevention of scale.

Capability funding is not the same as use case funding

A use case solves a specific problem in a specific context. A capability creates value across multiple use cases over time. Use case funding is variable and tied to a specific outcome. Capability funding is persistent and builds a compounding asset. Organisations that only fund use cases pay full cost for each initiative with no accumulating return. Those that fund capability reduce the marginal cost of every initiative that follows.

Intake must be structured and single

AI opportunities should enter through one mechanism, assessed against consistent criteria that include strategic alignment, value, risk, feasibility, and capability impact. When intake is fragmented, the portfolio reflects the negotiating strength of sponsors rather than the strategic priorities of the enterprise.

The ability to stop is a funding mechanism


Every initiative requires a clear outcome at defined review points: Scale, Reuse, Retire, or Kill. The capital released by stopping an underperforming initiative is not a saving. It is a reallocation to something with a better return. A portfolio that cannot stop cannot concentrate.

Reuse is an economic argument

Capabilities designed for reuse create compounding value across the portfolio. Each use case that draws on a shared platform costs less and deploys faster than one built from scratch. Reuse is not a technical preference. It is the logic of investment efficiency.

Outcomes are owned by the business

Technology enables AI capability. It does not realise business value. Value realisation requires process change, user adoption, and commercial accountability. Benefits ownership must sit with the business leader who controls those levers. When technology owns the outcome, it is measured as a technical result. When the business owns it, it is measured against the P&L.



THE AI PORTFOLIO GOVERNANCE FRAMEWORK

Operationalising these principles begins with controlling how opportunities enter and move through the portfolio. Without a structured mechanism, the portfolio is shaped by whoever applies the most pressure, not by what generates the greatest return on the capital allocated.




Figure 1: AI Portfolio Governance Funnel

The AI Portfolio Governance Funnel provides a single enterprise intake mechanism through which all AI opportunities pass. It applies progressively stricter criteria at each stage, filtering out initiatives that do not meet strategic, financial, capability, or feasibility thresholds before a funding decision is reached.

The five stages are sequential and cumulative. Strategic alignment ensures connection to enterprise priorities and a clearly defined business problem. Value and risk assessment evaluates financial impact, risk exposure, and regulatory fit. Capability impact and reuse assesses contribution to the broader capability portfolio, platform leverage, and the reusability of assets to be created. Data and feasibility readiness confirms that the data, integration, and technical conditions for delivery exist. Portfolio prioritisation ranks the initiative against others across the whole portfolio, with explicit trade-offs made at this stage, not earlier.

The output is a capital allocation decision: Fund, Defer, or Reject. The critical differentiator is the third stage. Most intake processes evaluate an initiative on its own merits. This funnel evaluates it on its contribution to the portfolio as a capital allocation to capability. That single addition changes the governance logic from local approval to collective investment management.



FUNDING, PRIORITISATION AND KILL CRITERIA

Once opportunities enter the portfolio, they must be prioritised and funded in a way that reflects their relative contribution to value, not the order of their arrival or the seniority of their sponsor.




Figure 2: AI Use Case Prioritisation Matrix

The AI Use Case Prioritisation Matrix maps initiatives across two dimensions: value and strategic impact on the vertical axis, and feasibility and data readiness on the horizontal axis. Four quadrants drive four different investment responses.

Scale Now initiatives are high-value and high-feasibility. These are the immediate priorities. They are ready for resourcing and scaling, and delay carries a direct opportunity cost. Strategic Bets are high-value but currently less feasible. They require investment in capability, data, or integration before they can scale. These are portfolio commitments, not near-term deliverables, and they need patient, persistent funding with a longer governance horizon. Optimise initiatives are low-value but highly feasible. They may deliver tactical efficiency gains, but investment is contained and time-limited. Reject initiatives are low-value and low-feasibility. They do not enter the portfolio, regardless of who is sponsoring them.

The matrix is most valuable for what it prevents. Organisations consistently over-invest in the Optimise quadrant because those initiatives are the easiest to execute and the easiest to report on. They under-invest in Strategic Bets because those initiatives require longer commitment and carry more uncertainty. A disciplined matrix surfaces that trade-off and forces it to be resolved at the right level.

Funding architecture matters as much as prioritisation. Capability funding is persistent. It is allocated to teams, platforms, and data infrastructure that enable multiple initiatives over time. Use case funding is variable. It is tied to specific outcomes within defined time boundaries. Organisations that do not make this distinction explicitly will repeatedly rebuild the same foundations across different business units, paying full cost each time with no compounding return.

At defined portfolio review points, every active initiative is evaluated against its progress and its continued right to capital.





Figure 3: AI Lifecycle Decision Model

The AI Lifecycle Decision Model applies a structured performance review across five dimensions: value realised, user adoption, risk profile, operational stability, and reusability of assets created. Four outcomes are possible.

Scale applies when value is high, adoption is strong, and the initiative is stable and repeatable. Investment increases and scope expands. Reuse applies when components or assets can be extended to new contexts, creating compounding value without full replication. Retire is the planned, orderly sunset of a capability that has been superseded or has served its purpose. It is a managed transition, not a failure. Kill is the active termination of an initiative that is not delivering value, is not being adopted, or carries risk that cannot be mitigated.

The distinction between Kill and Retire is commercially important and consistently misunderstood. Retirement is planned and applies to capabilities that have fulfilled their purpose. Kill is a governance decision to stop consuming capital without return. Capital released by a Kill decision must be explicitly reallocated to higher-priority investments. Without that reallocation discipline, stopping one initiative creates slack rather than strength.

“

Kill criteria are not failure management. They are how a portfolio concentrates capital on what works.




SCALING PATHWAYS AND REUSE

Scaling AI is not a continuation of a successful experiment. It is a categorically different challenge, and treating it as a natural extension of the pilot is one of the most consistent mistakes in enterprise AI programmes.

The pilot environment that produced a successful result was designed to remove friction: controlled scope, limited data complexity, a single business unit, a motivated sponsor with direct oversight. Enterprise scale reintroduces all of that friction, and adds more besides: fragmented data ownership, integration dependencies, inconsistent governance standards, and the full weight of change management across functions and geographies that the pilot never encountered.

The shift from experiment to scaled capability requires three conditions simultaneously. The technical assets must be designed for reuse, not just for the immediate problem. The operating model must absorb the capability into regular workflows, with clear business ownership and adoption accountability. And the governance model must transition from innovation oversight to operational performance management. Most organisations address the first. Very few address all three, which is why pilots that work technically still fail to generate business value.

Reuse is the most under-exploited lever in enterprise AI portfolios. When each business unit builds AI capabilities independently, the organisation pays full cost for every initiative with no compounding return. Data pipelines are rebuilt. Governance controls are duplicated. Model components are replicated across teams that are unaware of each other's work. The cumulative waste is significant and largely invisible, distributed across many budgets and many teams with no one accountable for the aggregate.



Platform investment converts that cost into portfolio leverage. When data infrastructure, model components, and governance mechanisms are designed as shared assets, the marginal cost of each subsequent initiative falls materially. The economics are straightforward: the second initiative costs less than the first. The fifth costs substantially less than the second. Over time, a well-designed platform base makes the portfolio self-funding in a way that a collection of standalone use cases never can.

Two directions of scale require different governance responses. Scaling depth means expanding a capability within its existing context: more users, more data, more transaction volume. This is primarily an operational and change management challenge. Scaling breadth means extending the capability to new contexts, new business units, or adjacent use cases. This requires a deliberate platform and reuse strategy. The distinction matters because organisations frequently attempt breadth scaling with a depth approach and achieve neither.



Platform investment is not a technology decision. It is a capital allocation decision. The portfolio governance layer must own it.

A related decision is build, buy, or partner. The same MIT GenAI Divide research found that AI initiatives delivered through external partnerships and vendor solutions reach production at roughly twice the rate of comparable internal builds. This is not an argument against internal capability. It is an argument for treating sourcing as a portfolio governance decision rather than a procurement one. The default assumption that strategic value requires internal construction is not supported by the evidence. For most use cases, the strategic question is not whether to build, but where building creates a defensible capability advantage and where buying or partnering accelerates value at lower risk. Portfolio governance is the right level at which to make that distinction, because only at portfolio level can the trade-offs between speed, control, cost, and capability accumulation be resolved coherently.

FAILURE MODES

Even with sound principles and the right governance structures, organisations fall into predictable traps. These are not abstract risks. They are patterns that are visible in most large enterprises today.

Funding the long tail instead of the high ground

The portfolio holds thirty initiatives. None has sufficient funding or executive attention to reach enterprise scale. Each has a sponsor, a status report, and a quarterly update. Collectively they consume the majority of available AI investment and deliver a fraction of the potential value. This is capital dilution presented as portfolio diversity. The decision to approve broadly rather than invest deeply is a governance failure at the most fundamental level. A portfolio that funds everything cannot concentrate on anything.

Treating continuation as a decision


An initiative that met its pilot criteria three quarters ago is still running. Adoption is flat. The business case has not been refreshed. No one has reviewed it formally, and no one has made the decision to stop it. Continuation is not a default position. It is a choice, and it is consuming capital that belongs to better investments. Organisations without explicit review cycles and kill criteria do not avoid this decision. They defer it until the funding runs out, with no learning captured and no capital released to something more productive.

Rebuilding what already exists

One business unit has spent six months building a document processing capability. Another is about to spend the same six months building its own version of the same capability. Neither team knows the other exists. No one at portfolio level has asked whether this capability already exists, could be extended, or should be built once and shared. The cost is not only the duplicated investment. It is the shared platform that was never built because no business unit was accountable for it and no governance mechanism required it.

Technology ownership of business outcomes

The AI capability works. The model performs well. The system is stable. And the value case made eighteen months ago has not materialised, because the process the capability was supposed to improve has not been redesigned, the users have not adopted it, and no business leader is being held accountable for the commercial result. Technology can build a capability. It cannot realise the value of one. When technology owns the benefit case, the benefit case measures technical performance. When the business owns it, it measures whether the investment has returned to the P&L.



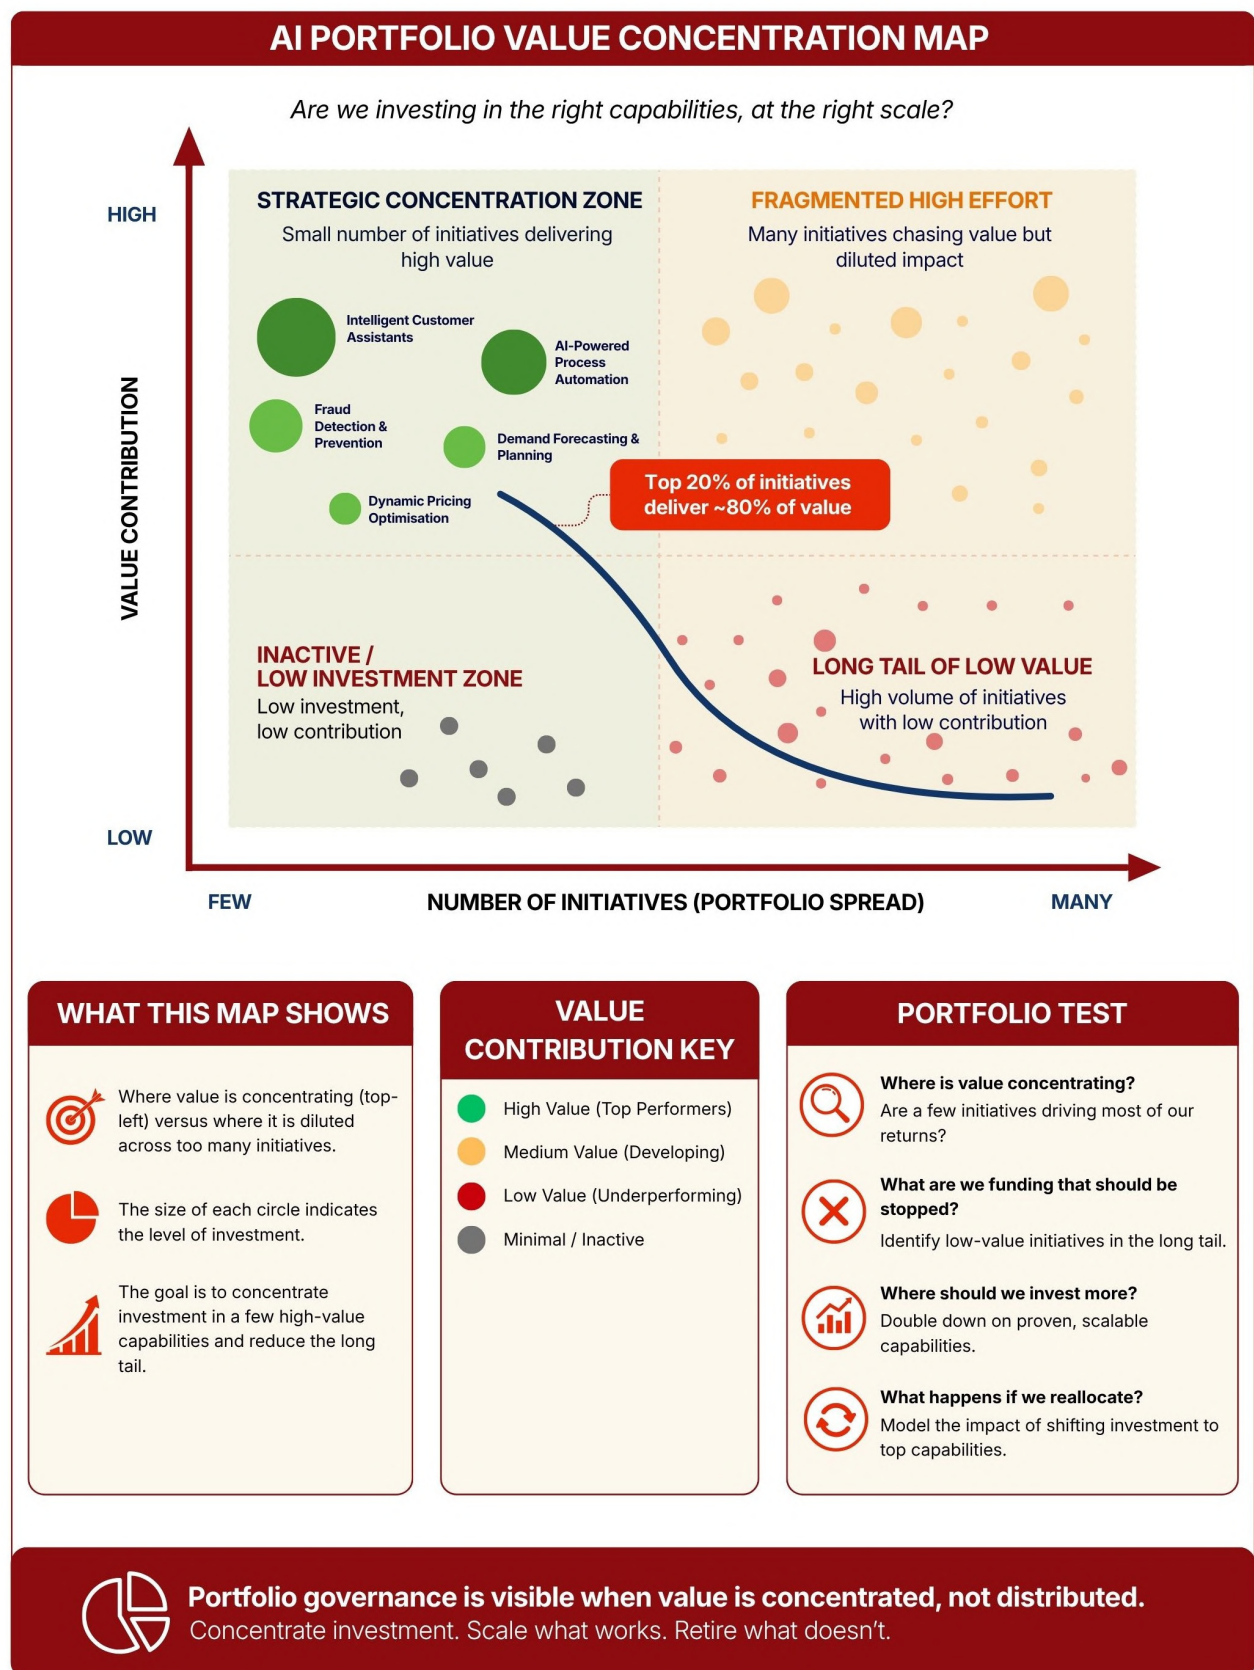



Figure 4: AI Portfolio Value Concentration Map

The AI Portfolio Value Concentration Map provides an immediate visual test of portfolio health. It plots all current initiatives against their value contribution, making the distribution of investment impact visible in a single view.

A healthy portfolio shows clear concentration: a small number of strategic capabilities delivering the substantial majority of value, with a deliberately contained and actively monitored long tail. When the map shows value distributed evenly across many initiatives, the portfolio is not functioning as a governance instrument. It is functioning as a catalogue, and the capital allocation conversation that should be happening at board level is not happening at all.

The map prompts the question that portfolio governance exists to answer: if we redirected the resources currently spread across the long tail into the top three capabilities, what would the impact be? An organisation that cannot answer that question, or chooses not to ask it, is not governing its AI portfolio. It is administering one.



WHAT LEADERS SHOULD DO NOW

Establish a single, structured intake mechanism

If AI opportunities are entering your portfolio through multiple channels, assessed on local criteria and approved by different decision-makers, you do not have a portfolio. You have a catalogue. Establish one intake mechanism, assessed against consistent criteria that include strategic alignment, value, risk, feasibility and capability impact. Every opportunity goes through it.

Apply portfolio logic to your current investments

Map your active AI initiatives against the prioritisation matrix. Ask which are Scale Now, which are Strategic Bets, which are Optimise and which belong in the Reject quadrant regardless of who sponsors them. The exercise will surface uncomfortable truths about where capital is concentrating and where it is not. That is the point.

Define kill criteria before funding decisions are made

Every initiative entering the portfolio should have a defined outcome at a defined review point: Scale, Reuse, Retire or Kill. Kill criteria established at funding are governance. Kill criteria applied after failure are damage control. Require them upfront.

Invest in capability, not only in use cases

Identify which of your AI investments are building reusable platforms, data infrastructure and shared model components, and which are paying full cost for a single-use outcome. Redirect a portion of use case funding into capability investment. The economics compound over time. Standalone use cases do not.

Assign business ownership of outcomes, not technology ownership of delivery

For each material AI initiative, confirm that a business leader is accountable for value realisation against the original business case. If technology holds the benefit case, it will be measured on technical performance. If the business holds it, it will be measured on commercial outcome. That distinction determines whether value actually lands in the P&L.

Report on value concentration, not activity

Replace activity-based portfolio reporting with a value concentration view. How much of your AI investment is generating measurable returns? Where is value concentrating? What proportion of active initiatives have a credible pathway to enterprise scale? If your current reporting cannot answer these questions, your governance model is measuring the wrong things.

BOARD AND EXECUTIVE QUESTIONS

For boards and executives, the shift to portfolio governance requires a different set of questions. Not more questions about individual initiatives, but sharper questions about the portfolio as a whole.

Where is our AI value concentrating today, and does that reflect a deliberate capital allocation decision or an accident of approval history?

Which initiatives are we actively choosing not to fund, and on what basis?

Are we investing in reusable capability, or only approving isolated use cases?

How many of our active initiatives have a credible, resourced pathway to enterprise scale?

What have we stopped in the last twelve months, and what did we do with the capital released?

Who owns the business outcomes of each major AI capability, and how are they being held to account?

Does our portfolio reporting tell us about value realised, or only about activity completed?

An organisation that cannot answer these questions clearly is not governing its AI portfolio. It is cataloguing it. The distinction is not semantic. It is the difference between compounding value and accumulating cost.

CLOSING INSIGHT

Most organisations do not have an AI problem. They have a governance problem.

The technology is capable. The use cases are real. The investment is committed. What is missing is the discipline to treat AI as what it actually is: a portfolio of capital allocations to capability, each of which must be governed, concentrated, scaled, or stopped with the same rigour applied to any other strategic investment.

The organisations pulling ahead are not running more experiments. They are making better investment decisions and holding themselves accountable for the outcomes. They build platforms, not just products. They stop what is not working and redeploy the capital. They require business leaders to own the value, not just the technology. And they ask, at every portfolio review, whether investment is concentrating where returns compound.

The organisations falling behind are not failing for lack of ideas or ambition. They are failing because their governance model was built for a world of bounded, sequential projects. It was not built to manage a compounding asset at enterprise scale, and it will not produce enterprise-scale returns until it is redesigned to do so.

“

AI portfolios are capital allocations to capability. The answer is not more pilots. The answer is knowing where to invest, having the discipline to concentrate, and the governance to scale what matters.



REFERENCES

- McKinsey & Company. The State of AI 2025: Adoption, Investment and Value Realisation. McKinsey & Company, 2025.
- McKinsey & Company. State of AI Trust 2026: Shifting to the Agentic Era. McKinsey & Company, 2026.
- Boston Consulting Group. Are You Generating Value from AI? The Widening Gap. BCG, 2025.
- Boston Consulting Group. AI Radar 2025: Closing the Impact Gap. BCG, 2025.
- Boston Consulting Group. Build for the Future 2025: AI Maturity and Value Creation. BCG, 2025.
- Boston Consulting Group. Enterprise as Code: An Operating Model for the AI Era. BCG, 2025.
- MIT Initiative on the Digital Economy (NANDA). The GenAI Divide: State of AI in Business. MIT, 2025.
- PwC. AI Business Predictions: Capturing Value in the Age of AI. PwC, 2025.
- Deloitte. AI ROI: The Paradox of Rising Investment and Elusive Returns. Deloitte, 2025.
- Deloitte. State of Generative AI in the Enterprise. Deloitte AI Institute, 2025.
- RAND Corporation. Understanding the High Failure Rate of AI Projects. RAND, 2024.
- S&P Global Market Intelligence. Enterprise AI Adoption and Abandonment Trends. S&P Global, 2025.
- Gartner. Forecast Analysis: Generative AI Project Outcomes. Gartner, 2024.
- National Institute of Standards and Technology. AI Risk Management Framework. NIST, 2023.
- International Organization for Standardization. ISO/IEC 42001 Artificial Intelligence Management Systems. ISO, 2023.
- European Commission. EU Artificial Intelligence Act. European Commission, 2024.
- Manoj Tavarajoo. Leading the AI Transformation. MyConsultancy, 2025.
- Manoj Tavarajoo. The AI Operating Model Playbook. MyConsultancy, 2026.

ABOUT THE AUTHOR

Manoj Tavarajoo has spent over two decades working with boards, executives and senior leaders across enterprise, digital and AI transformation. His work sits at the intersection of strategy, operating model design, governance and execution, where transformation ambition becomes measurable enterprise capability.

He is the author of *Leading the AI Transformation*, *The AI Operating Model Playbook* and *The AI Governance Imperative*, together forming a practical leadership trilogy on how organisations lead AI, organise for AI and govern AI at enterprise scale.

Manoj works through MyConsultancy, an independent advisory and transformation practice based in Australia, and is open to senior embedded leadership opportunities partnering with boards and executive teams navigating AI strategy, operating model design and governance at enterprise scale.




[@manojtavarajoo](https://www.linkedin.com/company/manojtavarajoo)

ABOUT MYCONSULTANCY

MyConsultancy works with boards and executive teams navigating the distance between AI ambition and operating reality. The practice specialises in AI strategy, operating model design, governance and enterprise transformation, helping organisations build the structures, decision-making disciplines and execution capability needed to scale AI responsibly.



www.myconsultancy.com.au



MyConsultancy