WHITE PAPER · PORTFOLIO & AI GOVERNANCE

Business Imagination Is the Scarce Asset Now

A portfolio-governance standard for turning cheap AI execution into measurable business value

Marco Policani
Enterprise Portfolio · PMO · AI Operating Governance
policani.net · linkedin.com/in/marcpolicani
July 2026

Executive premise

AI has made execution cheaper. It has not made value automatic.

That distinction is now the core leadership problem. Models can draft, summarize, classify, code, compare, search, translate, and synthesize at a speed and cost that would have looked implausible only a few years ago. At the same time, enterprise results remain uneven. McKinsey's 2025 global AI survey found that 88 percent of respondents report regular AI use in at least one business function, while only 39 percent report any enterprise-level EBIT impact from AI use.[1] McKinsey's 2026 agentic-organization discussion sharpened the same pattern: more than 80 percent of companies said they were not yet seeing bottom-line impact from their AI investments.[2]

The problem is not that AI cannot produce work. The problem is that most organizations are still asking it to do the work they already knew how to request.

That is why business imagination becomes scarce. In this paper, business imagination does not mean creativity as a soft trait. It means an operating capability: the ability to notice that a different workflow, decision, product, control, or customer interaction has become possible; test whether that possibility matters; compare it with other demands for capacity; and assign the right human, AI, data, governance, and funding assets to the work.

The practical question is no longer, "Can AI do this task?" Increasingly, it can. The better questions are: Which business problem is worth redesigning? What value would change if the work changed? What evidence would prove it? Which ideas deserve frontier exploration, which deserve cheap execution, which deserve deterministic automation, which deserve a human expert, and which should wait?

That is a portfolio question. It belongs in opportunity registers, business-value registers, funding gates, readiness reviews, and value-realization ledgers. It also belongs in how leaders evaluate AI ROI. AI ROI is not a post-hoc calculation attached to a pilot. It is the trace from a business hypothesis, through a redesigned workflow, into operational indicators, strategic outcomes, financial effects, adoption evidence, risk controls, and total cost of ownership.

Marco's public-safe product module proof-of-concept portfolio is relevant here because it makes that operating logic inspectable. The most relevant modules are not demonstrations of AI as production software. They are proof artifacts for how AI-shaped work can be routed: AI Opportunity Intelligence Review System, AI Artifact Lifecycle Governance System, Business Case System, Portfolio Prioritization Scoring Agent, Portfolio Capacity Sequencing Planner, Executive Portfolio Review Pack Builder, and Value Realization Governance Ledger.[3] Together, they show how ideas become evidence-bearing work rather than attractive activity.

The scarce asset is not the model. It is the operating imagination to choose better work, prove it, and resource it well.

The market signal: AI activity is broad; enterprise impact is not

The current evidence base is consistent on one point: AI activity has spread faster than enterprise value realization.

McKinsey's 2025 survey reported broad AI use, growing experimentation with AI agents, and a gap between use-case benefits and enterprise-level EBIT impact.[1] Its 2026 measurement framework makes the management implication explicit: many organizations cannot answer whether users are adopting AI as intended, whether it is paying off in customer experience or financial impact, or whether it is moving the business toward a more competitive model.[4] Gartner made the same failure mode operational in 2024, predicting that at least 30 percent of generative AI projects would be abandoned after proof of concept by the end of 2025 because of poor data quality, inadequate risk controls, escalating costs, or unclear business value.[5]

Those are not model-performance complaints. They are operating-design complaints.

The pattern matters because AI has made the easy measurement traps more tempting. Leaders can count users, prompts, tokens, documents generated, code suggestions accepted, meeting summaries created, chatbots launched, and pilots announced. Some of those measures are useful. None of them, by themselves, proves that the business became better.

McKinsey's five-layer AI measurement framework is useful because it separates technical performance, user adoption, operational KPIs, strategic outcomes, and financial impact.[4] That separation is exactly what an AI portfolio needs. A model can be technically healthy and still fail to change the workflow. Users can adopt a tool and still fail to improve cycle time, quality, risk, or customer experience. Operational KPIs can move locally and still fail to create material financial impact. ROI gets credible only when those layers connect.

The RAND Corporation's 2024 work on AI project failure makes the upstream problem even plainer. RAND found that misunderstandings and miscommunications about project intent and purpose cause more AI projects to fail than any other factor. Its recommendation is blunt: focus on the problem, not the technology.[6]

That is where business imagination enters. The imaginative act is not saying "use AI." It is recognizing the business problem clearly enough that AI can be considered alongside other options: redesign the workflow, change the decision rights, improve the data, build a module, buy a tool, automate a deterministic step, train a team, hire capacity, or wait.

What business imagination means operationally

Business imagination is the ability to create new option value from domain context.

That definition is intentionally practical. It does not ask leaders to be visionary in the abstract. It asks whether the people closest to the work can identify new possibilities that were previously too expensive, too slow, too manual, too fragmented, or too difficult to coordinate.

Three ingredients matter.

First, the organization needs people with real context. AI opportunity discovery is weak when it is conducted by people who know the model but not the workflow. The most useful questions often come from practitioners who know where handoffs fail, where status is reconstructed every week, where customers wait, where controls are duplicated, where exceptions hide, where analysts spend hours cleaning data before they can think, and where leadership decisions are delayed because the evidence pack is never stable.

Second, those people need working familiarity with AI capability. They do not need to become model engineers. They do need enough hands-on fluency to know what has become possible: large-scale comparison, first-pass synthesis, document structuring, classification, exception detection, controlled drafting, multi-source research, workflow scaffolding, and agentic handoff. Without touch, imagination stays generic.

Third, the organization needs permission and routing. A practitioner may see a high-value question, but if no one can spend a few hundred dollars of model time, create a sandbox, access source material, ask for review capacity, or register the idea without triggering a full project ceremony, the idea will die before it becomes visible.

This is why cheap execution and business imagination are complements, not opposites. Cheap models and routine automation are excellent for known work. They should be used aggressively where the task is clear, risk is bounded, quality can be checked, and the value is throughput. But the more commoditized known execution becomes, the more valuable the upstream question becomes: what should we ask AI to help us do now that we could not ask before?

There is an important boundary. Business imagination is not permission to chase novelty. It is disciplined option creation. It becomes useful only when routed through evidence, value, ownership, and review.

Why AI ROI starts before the pilot

Most AI ROI conversations start too late.

By the time a team is trying to prove ROI after a pilot, the most important decisions may already have been made badly: the wrong use case was selected, the workflow was never redesigned, the output was not tied to a decision, adoption was assumed, risk controls were afterthoughts, and total cost of ownership was not visible.

AI ROI should start as a value hypothesis in the idea register. The hypothesis does not need to be financially precise on day one, but it must be testable. A useful AI value hypothesis names:

- the business decision or workflow the idea is meant to improve;
- the current pain, cost, delay, quality gap, risk, or missed opportunity;
- the expected mechanism of value, such as work reduction, cycle-time compression, broader analytical coverage, better decision quality, risk reduction, customer improvement, or revenue enablement;
- the human owner of the outcome;
- the evidence required to move from exploration to reliance;
- the operating assets needed to test the idea;
- the reliance boundary that says what remains human-owned.

That makes ROI a chain, not a single number.

The chain can be measured in layers. Technical performance asks whether the system is reliable enough to use. Adoption asks whether the intended users actually incorporate it into the workflow. Operational KPIs ask whether the work changed: less rework, faster cycle time, fewer defects, better exception handling, lower cost per transaction, improved first-contact resolution, or higher throughput. Strategic outcomes ask whether the change matters to the business: customer experience, delivery reliability, compliance performance, retention, launch readiness, or portfolio decision speed. Financial impact asks whether the outcome produces lower cost to serve, revenue uplift, margin expansion, avoided exposure, or better capital allocation after costs are included.[4]

The research supports this layered view. Brynjolfsson, Li, and Raymond found in a 2025 Quarterly Journal of Economics study that a generative AI conversational assistant increased customer-support productivity by 15 percent on average, with larger gains for less experienced and lower-skilled workers and more mixed effects for the most experienced workers.[7] The point for portfolio leaders is not that every AI tool produces a 15 percent improvement. It is that credible measurement was tied to a specific workflow, a specific population, and observable work outcomes. That is what most enterprise AI business cases lack.

A useful ROI register should therefore include both direct and indirect value. Direct value includes hours avoided, cost-to-serve reduction, fewer defects, faster cycle time, lower escalation volume, and revenue lift. Indirect value includes faster decision cadence, improved evidence coverage, reduced control exposure, better onboarding, less context loss, and stronger adoption discipline. Indirect value should not be inflated into financial savings without evidence. It should be tracked as a leading indicator until enough proof exists to convert it into a financial claim.

The idea and business-value register

The idea register is where business imagination becomes governable.

Without a register, ideas arrive as anecdotes: a demo someone liked, a vendor claim, a hallway suggestion, a team experiment, a prompt that saved time, a script that "seems useful," a dashboard no one owns, or an agent prototype that starts shaping decisions before anyone reviewed its limits.

The register turns those inputs into comparable decision inventory.

A strong AI idea and business-value register should include at least these fields:

Field	Purpose
Idea or opportunity	Names the proposed change in plain business language.
Workflow affected	Prevents AI from floating above the work it is supposed to improve.
Decision supported	Names the decision, handoff, action, or customer outcome the work should improve.
Value mechanism	Work reduction, cycle time, quality, coverage, decision quality, risk control, revenue, or customer experience.
Current baseline	Captures what is known about current cost, delay, error, rework, risk, or friction.
Evidence required	Defines what must be proven before more capacity is assigned.
Data and source truth	Names the source material and whether it is reliable enough for the proposed use.
Reliance boundary	Clarifies what AI may draft, flag, classify, or recommend, and what stays human-owned.
Risk and control exposure	Identifies privacy, compliance, security, financial, reputational, and operational risk.
Asset fit	Compares human expertise, AI model, automation, vendor tool, data, process change, and governance capacity.
Recommended route	Explore, pilot, build, buy, automate, hire, wait, stop, or fold into existing work.
Owner and review cadence	Keeps accountability and follow-through visible.

This is not bureaucracy. It is how leaders prevent AI enthusiasm from becoming hidden work-in-progress.

The register should also distinguish between ideas and assets. An idea is a possible business change. An asset is something used to make that change real: a model, data set, script, dashboard, workflow module, vendor feature, review forum, policy, training path, subject-matter expert, or funding increment. Confusing the two is a common failure mode. Buying a tool does not mean selecting the right opportunity. Building a prototype does not mean the business case is sound. Assigning an analyst does not mean the workflow is ready to absorb the change.

Marco's AI Opportunity Intelligence Review System and AI Artifact Lifecycle Governance System are proof-of-concept modules for this distinction. The opportunity system asks whether the idea has a business purpose, value logic, owner, and review boundary. The artifact lifecycle system asks what kind of thing has been created, how mature it is, who owns it, what reliance it can support, and what review is needed before the business depends on it.[3]

That distinction is the difference between an AI demo and an AI operating system.

Selecting opportunities: the value test

Opportunity selection should compare AI ideas against all the other ways the organization can spend scarce attention.

The test is not whether AI can help. The test is whether AI is the right asset for the right work at the right time.

A practical selection model can use six gates.

Gate one: Is the problem worth solving? The organization should name the operating friction or business opportunity before discussing the model. If the problem is vague, the AI idea should stay in discovery.

Gate two: Is there a measurable value path? The value path can be direct or indirect, but it must be named. "Improve productivity" is too vague. "Reduce analyst time spent reconciling inconsistent portfolio fields before executive review" is a value path.

Gate three: Does the workflow need redesign? If AI is only accelerating a broken workflow, the organization may get faster waste. McKinsey's 2025 survey found that high performers were much more likely to redesign workflows and to define processes for human validation of model outputs.[1]

Gate four: Is the evidence good enough? The idea must identify source truth, baseline data, acceptance criteria, and the evidence needed to move from exploration to adoption.

Gate five: What is the risk of being wrong? Low-risk drafting and high-stakes financial, compliance, customer, or safety decisions should not share the same review model. NIST's AI Risk Management

Framework is useful here because it frames AI governance as a continuous cycle of governing, mapping, measuring, and managing risk across context, purpose, oversight, and lifecycle.[8]

Gate six: Which asset belongs to the work? Some work needs a frontier model to explore a new possibility. Some needs a cheaper model for repeatable execution. Some needs a deterministic script. Some needs better data. Some needs a workflow module. Some needs a human expert. Some needs a policy decision before any tooling.

This last gate is where portfolio discipline matters most. The organization should not assign its most expensive model, scarcest engineering capacity, or most senior human reviewers to every AI-flavored request. It should assign assets according to value, uncertainty, risk, repeatability, and evidence maturity.

Assigning assets to the right work

The phrase "AI resources" is too narrow. The useful resource pool includes at least seven asset classes.

Human assets include domain experts, process owners, reviewers, sponsors, risk partners, finance partners, operators, trainers, and change leaders. Their value is context, judgment, acceptance criteria, exception handling, and accountability.

AI assets include frontier models, lower-cost models, specialized models, retrieval systems, agents, copilots, and vendor features. Their value depends on matching capability to task and risk.

Data assets include source systems, documents, taxonomies, business rules, historical decisions, telemetry, case records, test data, and quality baselines. AI value is weak when source truth is weak.

Workflow assets include intake forms, decision trees, business cases, review packs, handoff templates, value ledgers, scoring rubrics, and operating cadences. These are often the missing layer between model capability and business impact.

Automation assets include scripts, validations, tests, integrations, deterministic checks, and reporting pipelines. These should handle repeatable rules where model judgment is unnecessary or risky.

Governance assets include review forums, escalation rules, reliance classifications, approval gates, risk controls, artifact lifecycle rules, and retirement paths.

Funding and capacity assets include budget, engineering time, analyst time, licenses, vendor support, training investment, and model spend.

The business-imagination discipline is to assign these assets deliberately. A high-uncertainty, high-upside opportunity may justify a short frontier-model exploration with senior domain review. A mature, repeatable workflow may justify moving execution to a cheaper model or deterministic

pipeline. A high-risk artifact may require a lifecycle classification and human signoff before any adoption. A promising idea with poor source truth may need data cleanup before AI is funded at all.

The decision is not "AI or no AI." The decision is asset fit.

The product-module proof of concept as inspectable operating evidence

The most relevant part of Marco's product-module proof-of-concept work is the route from idea to governed value.

The public portfolio should not be represented as production SaaS, employer deployment, autonomous decisioning, ML engineering, or measured enterprise AI savings. Its value is different: it makes the operating pattern inspectable. It shows how ambiguous AI and portfolio work can be decomposed into bounded modules with source truth, routing rules, reliance boundaries, review gates, output contracts, and decision-ready handoffs.

The relevant module chain is straightforward.

AI Opportunity Intelligence Review System is the intake gate. It asks whether an AI-shaped idea has a real workflow, a value mechanism, an owner, and a decision path.

AI Artifact Lifecycle Governance System is the reliance gate. It classifies the artifact as idea, draft, prototype, script, dashboard, agent, vendor feature, or operational artifact; then it names owner, evidence, allowed use, review status, update cadence, and retirement trigger.

Business Case System is the investment gate. It forces the idea to compete with other uses of capacity by stating options, assumptions, costs, risks, benefits, and the "do nothing" case.

Portfolio Prioritization Scoring Agent is the comparison gate. It helps evaluate value, urgency, risk, readiness, dependency, and confidence across competing opportunities.

Portfolio Capacity Sequencing Planner is the allocation gate. It asks when work can actually be absorbed, what capacity is constrained, what dependencies matter, and whether sequencing creates or reduces risk.

Executive Portfolio Review Pack Builder is the decision gate. It turns the opportunity into a reviewable evidence pack, not a loose narrative.

Value Realization Governance Ledger is the follow-through gate. It keeps the value claim visible after approval, through delivery, adoption, measurement, revision, and closure.

That chain is the practical answer to "where to select opportunities and assign assets to the right work." It does not start with the model. It starts with the business idea, then routes the work through evidence, decision, capacity, and realization.

A working operating model

The proposed operating model has five layers.

Layer one is the imagination register. This is where new business possibilities enter the system: workflow ideas, AI ideas, vendor claims, automation suggestions, analyst pain points, customer friction, risk signals, and executive questions. The register is lightweight at intake, but it must name the workflow, value mechanism, owner, and evidence gap.

Layer two is the value and risk screen. This screen determines whether the idea should be explored, stopped, delayed, folded into existing work, or converted into a business case. The screen should include value potential, risk of being wrong, source-truth readiness, workflow absorption, review burden, and asset fit.

Layer three is the asset-assignment decision. The portfolio decides whether the opportunity needs a human expert, frontier exploration, cheaper execution, deterministic automation, data remediation, vendor assessment, workflow redesign, training, policy review, or funding.

Layer four is the proof and scaling path. The work moves through exploration, pilot, MVP, initial scaling, or full scale only when evidence supports the next step. McKinsey's 2026 measurement article is useful here because it links value definition, attribution design, review cadence, decision gates, and evidence packs.[4]

Layer five is the value-realization ledger. This is where claims survive contact with reality. Expected value is tracked against actual adoption, operational change, strategic outcome, financial result, risk reduction, and total cost. If the evidence weakens, the claim is revised or retired.

The operating cadence should be simple:

- Weekly or biweekly intake review for new ideas and triage.
- Monthly opportunity review for experiments, pilots, evidence gaps, asset conflicts, and stop/go decisions.
- Quarterly portfolio review for funding, capacity, sequencing, realized value, and strategic fit.
- Periodic artifact lifecycle review for AI outputs that remain in use, especially dashboards, agents, scripts, summaries, decision aids, and vendor-enabled features.

The cadence creates the scarce capability: not more AI activity, but better judgment about where AI belongs.

What leaders should measure

The measurement set should be small enough to use and complete enough to prevent self-deception.

For each AI opportunity, measure the following:

Measurement layer	Example measures	Owner
Technical performance	Accuracy, latency, hallucination or error rate, token cost, drift, system availability.	Technical or tool owner.
Adoption and trust	Eligible users, active use, workflow penetration, acceptance rate, override rate, reviewer edits, user confidence.	Product, process, or operations owner.
Operational impact	Cycle time, rework, defect rate, throughput, backlog age, exception detection, escalation volume, cost per case.	Process owner.
Strategic outcome	Decision speed, customer experience, compliance performance, delivery reliability, portfolio signal quality, revenue enablement.	Business owner.
Financial impact	Cost reduction, cost avoidance, revenue uplift, margin impact, avoided exposure, total cost of ownership.	Finance and sponsor.
Governance health	Reliance boundary, review completion, evidence freshness, incident log, retirement trigger, risk acceptance status.	Governance or risk owner.

The governance-health layer is often missing. It matters because AI can create plausible outputs that outlive their evidence. A summary, dashboard, model-assisted analysis, or prototype agent can become part of business rhythm without anyone remembering when it was last validated. That is how experiments become shadow operating systems.

Measurement should also preserve uncertainty. Not every benefit can be financially attributed immediately. That is acceptable if the register labels evidence quality. A cycle-time reduction with clean baseline data is stronger than a manager's estimate. A controlled comparison is stronger than before-and-after anecdotes. A leading indicator is not a realized financial benefit. A risk-reduction claim should name the exposure reduced and the evidence for reduction. A value claim that cannot be verified should not be promoted to ROI.

The paper's core position is not that every AI use case needs a heavyweight financial model. It is that every AI use case needs a clear value standard appropriate to its cost, risk, and reliance level.

Claims this paper does not make

This paper does not claim that AI ROI is impossible. It claims that ROI is often measured too late, too narrowly, or too far away from the workflow.

It does not claim that cheap models are bad. Cheap execution is a gift when the work is known, repeatable, low-risk, and testable. The claim is that cheap execution does not answer which work should exist.

It does not claim that every idea needs a product team. Some ideas should remain prompts, templates, scripts, or human practices. The governance point is to decide that deliberately.

It does not claim that Marco's proof-of-concept modules are production systems or employer deployments. They are public-safe operating artifacts that make a pattern visible: opportunity review, artifact lifecycle governance, business-case routing, portfolio scoring, capacity sequencing, executive review, and value realization.

It does not claim measured AI cost savings from the product-module portfolio. The claim is about operating design, not quantified enterprise savings.

The decision standard

Business imagination is now scarce because execution is increasingly abundant.

When everyone can summarize the same meeting, draft the same document, generate the same code pattern, build the same chatbot, and route the same known task to a cheaper model, advantage moves upstream and downstream. Upstream, advantage comes from better questions: which workflow, decision, customer interaction, control, or product experience should now be redesigned? Downstream, advantage comes from proof: did the new work change cost, cycle time, quality, risk, adoption, decision speed, customer experience, or financial outcome?

The operating standard is simple:

AI opportunities should not compete on novelty. They should compete on value logic, evidence readiness, workflow fit, risk control, asset fit, capacity realism, and measurable follow-through.

That is what an idea register is for. That is what a business-value register is for. That is what a portfolio review is for. That is what a value-realization ledger is for. They make imagination accountable without killing it.

The leaders who win with AI will not be the ones who ask models to do the old work slightly faster. They will be the ones who build an operating system for discovering better work, proving it, assigning assets to it, and stopping what does not earn the next increment of trust.

Notes and sources

1. Alex Singla, Alexander Sukharevsky, Bryce Hall, Lareina Yee, Michael Chui, and Tara Balakrishnan, "The state of AI in 2025: Agents, innovation, and transformation," McKinsey & Company, November 5, 2025. Verified July 5, 2026. https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai

2. Alexis Krivkovich and Lucia Rahilly, "AI is everywhere. The agentic organization isn't-yet," The McKinsey Podcast, April 2, 2026. Verified July 5, 2026. https://www.mckinsey.com/capabilities/people-and-organizational-performance/our-insights/ai-is-everywhere-the-agentic-organization-isnt-yet

3. Marco Policani, public-safe AI and governance module portfolio catalog and related module pages, updated June 9, 2026; local evidence cross-checked against `E:\Codex\job-search\Resume Assets Engine\case-ai-practice.md` on July 5, 2026. Portfolio entry point: https://policani.net/modules.html

4. Johannes-Tobias Lorenz, Joshan Cherian Abraham, Robert Levin, and Douglas Ziman, "From promise to impact: How companies can measure-and realize-the full value of AI," McKinsey & Company, April 24, 2026. Verified July 5, 2026. https://www.mckinsey.com/capabilities/quantumblack/our-insights/from-promise-to-impact-how-companies-can-measure-and-realize-the-full-value-of-ai

5. Gartner, "Gartner Predicts 30% of Generative AI Projects Will Be Abandoned After Proof of Concept By End of 2025," July 29, 2024. Verified July 5, 2026. https://www.gartner.com/en/newsroom/press-releases/2024-07-29-gartner-predicts-30-percent-of-generative-ai-projects-will-be-abandoned-after-proof-of-concept-by-end-of-2025

6. Benjamin Boudreaux, Bruce J. McConnell, Jonah Czerwinski, and David Luckey, "The Root Causes of Failure for Artificial Intelligence Projects and How They Can Succeed: Avoiding the Anti-Patterns of AI," RAND Corporation, 2024. Verified July 5, 2026. https://www.rand.org/pubs/research_reports/RRA2680-1.html

7. Erik Brynjolfsson, Danielle Li, and Lindsey Raymond, "Generative AI at Work," The Quarterly Journal of Economics, Volume 140, Issue 2, May 2025, pages 889-942. Verified July 5, 2026. https://academic.oup.com/qje/article/140/2/889/7990658

8. National Institute of Standards and Technology, "AI Risk Management Framework Core," excerpt from AI RMF 1.0, 2023. Verified July 5, 2026. https://airc.nist.gov/airmf-resources/airmf/5-sec-core/

About the author
Marco Policani is an enterprise portfolio, PMO, and AI operating-governance leader. He builds portfolio governance systems where AI carries the analytical load and named humans own every decision.

Portfolio and governance library: policani.net/governance · Full portfolio: policani.net · LinkedIn: linkedin.com/in/marcpolicani