

The Ratio Is 7:1. Most Enterprises Invert It.

Whitepaper

Bart de Graaff · BdG Advisory · October 2026

Enterprises build agents, models and automation first, then wonder why nothing scales. The 7:1 ratio is a design heuristic, not a statistic, and it describes where the work actually is: seven questions of authority and control for every unit of technology. Invert it and governance still gets built, only later, under pressure and at a higher price.

Seven parts governance. One part AI.

Most enterprises build AI in the wrong order. They buy the capability first and look for the authority to use it later.

The ratio should be 7:1. Seven parts of an AI deployment should be governance, controls and authority structures. One part should be new execution capability: the models, agents and automation. Most enterprises invert it. They put most of their money, talent and senior attention into the one part, and treat the seven as a compliance step at the end. When the ratio inverts, every scaling attempt stalls. Not because the AI failed, but because the authority model was never there.

To be precise about the number: 7:1 is a design heuristic. It is mine, drawn from running businesses and watching technology programmes scale or stall. No study measured it, and this paper does not pretend one did. It is a statement about where effort, budget and leadership attention must go before AI touches a consequential decision. The evidence in what follows shows the symptoms of inversion and the direction that research, standards and the better-run deployments are moving. It does not prove the number. The number is there to force a conversation that most AI budgets avoid.

This paper defines the seven parts, explains why it is rational for enterprises to invert the ratio, shows what inversion looks like in production, makes the case for the right ratio, takes on the strongest objection, and ends with the moves that restore the balance, in order.

What the seven parts are

"Governance" is too vague a word to budget against. So make it concrete. The seven parts are seven questions. Each must have an owner and an answer before an AI system acts on a consequential decision.

One: the decision inventory. Which decisions will AI touch, and which of them are consequential? Most organisations can list their applications. Very few can list their decisions. Without the list, nobody can say where AI is safe to act, where it should only recommend, and where it should stay out.

Two: decision rights. Who owns each decision, and its outcome, once a machine shapes it? When the model drafts and a person approves, the accountability map usually still points at the person. The authority has moved. Decision rights make the new arrangement explicit.

If you do not assign them, a court or regulator will. When an airline argued before a Canadian tribunal in 2024 that its website chatbot was a separate legal entity responsible for its own actions, the tribunal member called it a remarkable submission and held that the airline was responsible for all the information on its website. The deploying organisation owns the outcome whether or not it designed for that. Part two is deciding who inside the organisation carries it.

Three: data authority. Who decides which data may reach which model, for what purpose? Models perform better with more context, which creates constant pressure to widen access. Someone has to own the limit.

The failure is usually gradual. A marketing team connects a model to customer records to personalise campaigns. A service team, seeing it work, points the same connection at customer queries. Within months the model is answering complaints with purchase history it was never approved to use for that purpose. Nobody made a bad decision. Nobody made a decision at all.

Four: identity and access for machines. An agent with system credentials is a new kind of actor in the enterprise. It needs its own identity, its own permissions and the same least-privilege discipline as a new employee. IBM's 2025 study of 600 organisations that had suffered a data breach found that 13% reported breaches of AI models or applications, and 97% of those lacked proper AI access controls.

Five: entrustment levels. Is this system recommending, executing under supervision, or executing alone? What evidence moves it up a level, and what moves it down? In *AI Hallucinations Are Emotions Without a Therapist* I argue that this graduated autonomy is the core containment structure for systems that are confident and sometimes wrong.

Six: evaluation and evidence. How do you know the system works, in this workflow, against the outcomes the business cares about? Evaluation before production, monitoring after it, and an evidence trail that lets someone reconstruct why a decision was made.

Seven: containment and recovery. What happens when it goes wrong? Who can stop it, how fast, and how do you roll back? The EU AI Act requires that people overseeing high-risk systems can interrupt the system through a "stop" button or a similar procedure. Most enterprise deployments have no equivalent. Containment also covers the incident playbook and the contract terms that allocate liability.

The one part: execution capability. The model, the agent, the integration, the automation. It matters. It is also the part that is easiest to buy.

Laid out like this, the number stops looking arbitrary. Each of the seven questions is a piece of design work in its own right, with its own owner, its own negotiation and its own upkeep. Each is roughly as demanding as getting the capability to work. Seven questions, each about as heavy as the technology: that is the shape of the heuristic. The precise ratio matters less than the shape.

Ease of purchase is the other reason the ratio is not 1:1. Capability is increasingly available from many suppliers, and the supply is noisy rather than scarce. Gartner estimates that only about 130 of the thousands of agentic AI vendors are real; the rest are rebranding existing products. The seven parts, by contrast, are specific to you. No vendor can tell you who owns your credit decisions, which customer data may reach which model, or when your claims agent should hand a case to a person. What can be bought should take a small share of internal effort. What only you can build should take most of it.

Engineers saw a version of this a decade ago. Google's 2015 paper on technical debt in machine learning showed that only a small fraction of real-world ML systems is composed of the ML code; the surrounding infrastructure is vast and complex. That was about data pipelines and monitoring. Now that AI systems act, the surrounding system has to include authority as well.

Why enterprises invert it

Nobody decides to invert the ratio. It inverts because every actor involved is behaving rationally.

Vendors sell the one part. A demonstration is capability. Nobody demonstrates a decision-rights matrix. The procurement conversation is therefore about what the technology can do, and the seven parts are left as "implementation considerations" for the buyer.

Boards and executive teams ask what the organisation is doing with AI. The visible answer is capability: pilots launched, agents deployed, partnerships announced. A decision inventory is not an announcement. So capability is what gets reported, and what gets reported gets funded.

Budgets reward a launch date. Capability projects have one. The seven parts look like overhead with no launch, sitting in a cost centre with no demo. In an annual planning round they lose.

Risk, legal and security are brought in late, as approvers rather than designers. They arrive at a working pilot with a list of objections. Governance therefore shows up as friction, which confirms the belief that governance slows things down. The cycle reinforces itself.

Finally, the seven parts are politically expensive and the one part is not. Buying an agent requires nobody to give anything up. Deciding who owns a decision that used to sit across three functions redistributes authority, and somebody loses. In *Why AI Transformation Fails* I argue that this redistribution is exactly where organisations push back. The inverted ratio is the path of least political resistance.

There is a measurement trap underneath all of this. Capability is easy to count: models licensed, agents deployed, pilots launched, hours saved in a demo. The seven parts produce value mostly as things that did not happen: the breach that was contained, the wrong decision that was escalated, the rollback that took minutes instead of weeks. Organisations fund what they can count. A dashboard full of capability metrics and empty of control metrics is the inverted ratio, made visible.

The result is familiar. The pilot works, because it runs in protected space with a dedicated team, a narrow scope and nothing to integrate. Scaling meets the missing seven parts all at once.

The data shows the stall. In McKinsey's 2025 global survey, out of 25 attributes tested, the redesign of workflows had the biggest effect on whether organisations saw bottom-line impact from generative AI. Only 21% of organisations using it had fundamentally redesigned even some workflows. S&P Global Market Intelligence found that 42% of companies abandoned most of their AI initiatives in 2025, up from 17% the year before, and that the average organisation scrapped 46% of its proofs of concept before production, citing cost, data privacy and security risks.

Gartner predicts that over 40% of agentic AI projects will be cancelled by the end of 2027, because of escalating costs, unclear business value or inadequate risk controls. Read those three reasons again. Unclear value is a missing decision inventory. Inadequate risk controls are missing parts four to seven. Escalating cost is what happens when capability is bought before anyone knows which decisions it is for.

What inversion looks like in production

The inverted ratio has a signature. Capability is deployed. Something goes wrong. The organisation then builds the missing parts, under pressure, in public, at a higher price.

In July 2025, a software founder was running a public experiment building an application with an AI coding agent. He had told it explicitly not to make changes without approval during a code freeze. The agent ran destructive commands anyway and deleted a production database holding records on 1,206 executives and 1,196 companies. The platform's chief executive called it unacceptable, and the company shipped fixes: automatic separation of development and production databases, a planning-only mode that prevents unapproved changes, and one-click restoration from backup.

Look at that list. Separation of environments is part four. A planning-only mode is part five. One-click restore is part seven. None of the fixes was a better model. The control had lived

in a prompt, which is to say it did not exist. The seven parts were built after the incident instead of before it.

In July 2025, Commonwealth Bank of Australia announced that 45 customer service roles would be replaced by an AI voice bot. In August, after pressure from the Finance Sector Union, it reversed the decision, saying the roles were not redundant and that it "should have been more thorough in our assessment". The capability worked well enough to deploy. What was missing was the evidence: a measured understanding of what the decision to remove the roles actually depended on. That is part six.

In May 2025, Klarna's chief executive explained why the company was investing in human customer service again after an aggressive shift to AI. Cost had been "a too predominant evaluation factor", he said, and the result was lower quality. The capability performed against the metric it was given. The metric was the problem. Evaluation that measures the wrong outcome is part six done badly, and it is as dangerous as part six not done at all.

Then there is the inversion nobody approves. When the sanctioned path to AI is slow or absent, people build their own one part. IBM found that one in five organisations studied reported a breach involving shadow AI, and those with high levels of it faced about US\$670,000 more in breach costs than those with little or none. And 63% of the breached organisations had no AI governance policy or were still developing one. Shadow AI is the inverted ratio at its purest: capability with zero governance, procured by whoever needed it that week.

The pattern across all four is the same. Governance is never skipped. It is deferred to the incident review.

Deferral is expensive for a structural reason. Once an AI system has been running in a workflow for a few months, people have built habits, handoffs and targets around its outputs. Adding decision rights, access limits or a supervision level at that point means retraining staff, renegotiating contracts, re-approving processes and, often, explaining to customers why the service has changed. The seven parts cost far less when they are designed into the workflow than when they are retrofitted onto it.

What the right ratio looks like

The right ratio looks slower at the start and faster everywhere after.

A global wealth manager provides the clearest public example, with a caveat: the account comes from its AI vendor's own enterprise guide. Morgan Stanley's answer to doubts about AI in a relationship business was to conduct intensive evaluations for every proposed application before it reached production, including human trainers grading model output against responses from expert advisers. Those evaluations gave the firm the confidence to roll use cases into production. According to the same account, 98% of its advisers now use the tool daily, and access to documents rose from 20% to 80%. Part six came first. Adoption followed.

Independent research points the same way from a different angle. BCG's 2024 study of AI value found that 74% of companies had yet to show tangible value from AI. The leaders followed a different allocation: 10% of effort on algorithms, 20% on technology and data, and 70% on people and processes. That is BCG's heuristic, not mine, and it slices the problem differently: its 70% includes change management and skills, where 7:1 is about authority and control. But both put the algorithm at roughly a tenth of the work. Both say that the part most organisations fund first is the smallest part of the job.

McKinsey's survey found that a chief executive's oversight of AI governance was one of the elements most correlated with higher self-reported bottom-line impact from generative AI, particularly at larger companies. The standards bodies agree on structure. NIST's AI Risk Management Framework treats governance as a cross-cutting function infused throughout mapping, measuring and managing risk, not as a fourth step at the end.

Regulation is moving in the same direction. For high-risk systems, Article 14 of the EU AI Act requires that the people overseeing them can remain aware of automation bias, override or reverse the output, and interrupt the system so that it comes to a halt in a safe state. Every one of those is a governance design requirement. None of them is a capability requirement. Where the law already reaches, the seven parts are not optional. Where it does not yet reach, it is a fair guide to where it is heading.

Consider how this plays out as a pattern. A claims operation wants an AI agent to triage incoming claims. A capability-first team would buy the agent and start a pilot. A governance-first team starts with the decision inventory and finds that triage is really a dozen decisions, of which a few are low-risk, a few are consequential, and a few touch regulatory obligations. It assigns an owner to each. It gives the agent its own identity, read access to the claims system and no write access. It sets the agent to recommend on the consequential decisions and act only on the low-risk ones, with written criteria for promotion and demotion. It builds the evaluation set from real past claims. Only then does it choose the model. If the evaluation shows that a cheaper model is good enough, it buys the cheaper model.

The first deployment takes longer to demonstrate. The second takes much less time, because the identity model, the entrustment framework, the evaluation harness and the incident playbook already exist. That is the compounding effect of the right ratio. The seven parts are infrastructure. They are built once and reused. Capability is specific to each use case.

This also answers the question of whether the ratio holds forever. The build cost of the seven parts falls as they are reused. The run cost does not. Monitoring, re-evaluation, entrustment reviews, access recertification and incident response are permanent. Across the portfolio, the ratio holds.

You can read your own ratio without a consultant. Take the AI programme's last budget, its staffing plan and the agendas of its last six steering meetings. Mark every line as either capability (models, agents, platforms, integration, the engineers building them) or one of the seven parts (decision mapping, ownership, data authority, machine identity, entrustment,

evaluation, containment). If capability dominates all three, the ratio is inverted, whatever the strategy document says. If the seven parts appear only as a line called "governance" with one person against it, it is inverted badly.

The test is crude by design. It does not need to be precise to be useful. If the result is uncomfortable, that is the point: it changes the next budget conversation.

The strongest objection: governance-first is how you lose

The best counter-argument deserves a fair hearing, because parts of it are right.

Speed wins. Capability compounds through use: teams learn what works by deploying, not by designing. Competitors who ship will learn faster than competitors who govern. Governance is also how incumbents protect themselves; in the perspective on structural transformation I make that argument myself. "Seven parts governance" sounds like seven committees. And shadow AI proves that over-governed official channels push people towards ungoverned tools. Finally, 7:1 is a heuristic, not a measurement. Why should anyone reorganise a budget around it?

Five answers.

First, the seven parts are not approval layers. A committee is what organisations build when they lack structure. Decision rights, scoped machine identities, evaluation gates and rollback are what let you say yes quickly. When the controls are built into the system, low-risk decisions move without a meeting. Autonomy can only be granted where it can be contained. Without the seven parts, the only safe answer to every new use case is another review.

Second, speed to demo is not speed to value. Scrapping 46% of proofs of concept before production, as the average organisation in the S&P survey did, is not fast learning. It is spending twice.

Third, shadow AI supports the argument rather than undermining it. People route around the official path when the official path is slow or missing. The answer is not less governance. It is a sanctioned path with the seven parts already built in, so that it is faster than the shadow path. IBM's numbers show what the alternative costs.

Fourth, the objection is right that legacy governance cannot do this work. The existing risk committee, asked to govern AI, will do what it was built to do: protect the existing structure. That is why the seven parts must be designed as new operating structure with new owners, not delegated to the committee that already exists.

Fifth, on the number itself. It is a heuristic and should be held as one. Exploration in a sandbox can and should run inverted: a team testing what a model can do needs capability, not an entrustment framework. The ratio applies at the point where AI touches consequential decisions, which is the point where most programmes stall. Its value is diagnostic. If your AI budget, hiring plan and steering-committee agenda are mostly about capability, you have a plan to buy what you cannot absorb.

Restoring the ratio: what leaders change, in order

The sequence matters. Each move makes the next one possible.

First, inventory before you buy. Pause new capability spend until every live pilot has a decision inventory: which decisions it touches, which are consequential, and which it should never make. This replaces the use-case backlog as the programme's master list. A use case is something you want to build. A decision is something you are accountable for.

Second, name the decision owners. For each consequential decision AI will shape, name one accountable person, and redraw the accountability map so it reflects where authority actually sits.

Third, give machines identities. Every agent gets its own credentials and the least access it needs. Separate experimentation from production. Do not let any control live only in a prompt.

Fourth, set entrustment levels. Recommending, supervised or autonomous, with written criteria for promotion and demotion based on observed performance in the workflow.

Fifth, make evaluation the production gate. Nothing reaches a consequential decision without an evaluation against the outcomes the business cares about, not just cost or accuracy. Build the evaluation sets from your own history, graded by your own experts. Keep evaluating after go-live, because the workflow, the data and the model will all change underneath the system.

Sixth, build containment before go-live. Stop mechanisms, rollback, incident playbooks and contract terms that allocate liability, tested before the first real customer sees the system.

Seventh, fund the capability. Now buy or build the one part, choosing it against the evaluation rather than the demo. Then track the ratio: in the budget, the hiring plan and the agenda of whoever governs the programme.

And stop four things. Stop counting pilots as progress. Stop putting controls in prompts. Stop inviting risk, legal and security to approve what they should have helped design. Stop letting vendor demonstrations set the roadmap.

One test tells you whether the moves have landed. Pick any AI system that touches a consequential decision and ask whoever owns the programme seven questions. Which decisions does it touch? Who owns them? What data may it use? What can it access? What is it allowed to do alone? How do you know it works? How do you stop it? If the answers come quickly, from named people, with evidence, the ratio is right. If the answers start with "the vendor" or "we are working on it", it is not.

Anyone can buy the one part. Only you can build the seven.

What to do about it

- **Inventory before you buy.** Pause new capability spend until every live pilot has a decision inventory: which decisions it touches, which are consequential and which it must never make.
- **Name the decision owners.** Assign one accountable person to each consequential AI-shaped decision and redraw the accountability map to match where authority actually sits.
- **Give machines identities.** Issue every agent its own credentials with least-privilege access, separate experimentation from production, and keep no control only in a prompt.
- **Set entrustment levels.** Classify each system as recommending, supervised or autonomous, with written promotion and demotion criteria based on observed performance.
- **Gate production on evaluation and containment.** Require evaluation against real business outcomes, plus tested stop, rollback and incident procedures, before any system reaches a consequential decision.
- **Then fund capability, and track the ratio.** Choose the model against the evaluation rather than the demo, and track the balance in the budget, the hiring plan and the governance agenda.

Sources

- IBM, Cost of a Data Breach Report 2025 press release: 13% of organisations reported breaches of AI models or applications, 2025: <https://newsroom.ibm.com/2025-07-30-ibm-report-13-of-organizations-reported-breaches-of-ai-models-or-applications,-97-of-which-reported-lacking-proper-ai-access-controls>
- Yahoo News Canada / The Hill, Air Canada must pay refund promised by AI chatbot, tribunal rules, 2024: <https://ca.news.yahoo.com/air-canada-must-pay-refund-040527669.html>
- Regulation (EU) 2024/1689 (EU AI Act), Article 14: Human oversight: <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>
- Gartner, Gartner predicts over 40% of agentic AI projects will be canceled by end of 2027, 2025: <https://www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027>
- Sculley et al. (Google), Hidden Technical Debt in Machine Learning Systems, NeurIPS 2015: <https://proceedings.neurips.cc/paper/2015/file/86df7dcfd896fcdf2674f757a2463eba-Paper.pdf>
- McKinsey, The state of AI: How organizations are rewiring to capture value, 2025: <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-how-organizations-are-rewiring-to-capture-value>

- CIO Dive, S&P Global Market Intelligence survey on abandoned AI initiatives, 2025:
<https://www.ciodive.com/news/AI-project-fail-data-SPGlobal/742590/>
- eWeek, AI agent wipes production database, then lies about it, 2025:
<https://www.eweek.com/news/replit-ai-coding-assistant-failure/>
- Information Age (ACS), CBA reverses AI-driven job cuts, admits error, 2025:
<https://ia.acs.org.au/article/2025/cba-reverses-ai-driven-job-cuts--admits--error-.html>
- Fortune, Klarna CEO on reinvesting in human customer service, 2025:
<https://fortune.com/2025/05/09/klarna-ai-humans-return-on-investment/>
- OpenAI, AI in the Enterprise (Morgan Stanley case, vendor account), 2025:
<https://cdn.openai.com/business-guides-and-resources/ai-in-the-enterprise.pdf>
- BCG, AI Adoption in 2024: 74% of companies struggle to achieve and scale value, 2024: <https://www.bcg.com/press/24october2024-ai-adoption-in-2024-74-of-companies-struggle-to-achieve-and-scale-value>
- NIST, AI 100-1: Artificial Intelligence Risk Management Framework (AI RMF 1.0), 2023:
<https://doi.org/10.6028/NIST.AI.100-1>

Read online: bdgadvisory.com/perspectives/enterprise-ai-deployment