

Artificial Intelligence

How AI Governance Defaults Shape Organizational Learning

Stijn Viaene




Image Credit | wacomka

Shift governance effort from building approval layers and processes toward building shared infrastructure, clear conventions, and routines that help learning circulate.

In response to generative AI's rapid diffusion across society and in professional work, organizations have developed AI strategies, created chief AI officer¹ roles and launched AI centers of excellence. In addition, they are establishing AI governance structures and mechanisms at a remarkable pace. All these arrangements aim to coordinate adoption, manage risk, and accelerate organizational learning.

RELATED ARTICLES

David J. Teece, Margaret A. Peteraf, and Sohvi Leih, "[Dynamic Capabilities and Organizational Agility: Risk, Uncertainty, and Strategy in the Innovation Economy](#)," California Management Review 58, no. 4 (2016): 13–35.

Christian Sciuk, Simon Engert, Maren Gierlich-Joas, and Thomas Hess, "[How Companies Navigate the \(Un\)Chartered Waters of Digital Transformation](#)," California Management Review 68, no. 1 (2025): 5–31.

RELATED TOPICS

[Digital Transformation](#)

[Organizational Design](#)

[Strategic Agility](#)

[Capabilities](#)

But as AI governance takes shape, a fundamental choice is emerging. Organizations are defaulting to one of two principles: approval-first or enablement-first governance, and that choice quietly shapes how organizations learn about AI in practice. While both approaches

may include elements of the other, the default, meaning what happens when situations are ambiguous, shapes everything that follows: culture, learning velocity, innovation outcomes, and how effectively organizations adapt to rapidly evolving AI capabilities.

Three factors make this choice particularly consequential now.

First, generative AI is accessible in ways previous enterprise technologies were not.

Anyone can use tools like ChatGPT, Gemini, or Claude through natural language conversations, often outside formal deployment cycles and sometimes beyond direct organizational visibility. No specialized training is required. Access is no longer limited to staged rollouts to designated user groups. This means experimentation can emerge across the workforce simultaneously, as employees identify and develop use cases in their day-to-day work². Traditional approval processes were designed for controlled technology rollouts to trained users. They struggle to coordinate experimentation occurring at this distributed scale and speed.

Second, generative AI capabilities evolve faster than organizations can fully stabilize their plans and governance arrangements around them.

Governance policies designed for text-based chatbot use may not anticipate multimodal, agentic, or code-generating capabilities that arrive months later. Capabilities and applications are evolving rapidly, making it difficult to predict how even the near-term future will unfold.³ This means organizations must learn what works through experimentation in practice. Traditional governance regimes assume relatively stable technology that allows for careful planning and staged rollouts. With generative AI, that stability is often temporary. Organizations that slow down to plan carefully risk finding that the learning opportunity has already moved on by the time implementation begins.

Third, governance structures determine what becomes visible and what goes underground.

When organizations establish governance, they make choices about whether experimentation surfaces or disappears. Get the default wrong, and valuable work may become invisible to the structures meant to coordinate it. This matters especially with accessible technology that people can adopt through consumer tools outside organizational systems. The organization may ban the use of generative AI tools, but that

does not mean that employees don't have access to their power on the smartphones they bring into the office, placing organizational learning partially outside formal governance visibility.⁴

This article draws on observations from my work with executives in organizations establishing AI governance, primarily through executive education programs and advisory engagements. I observed these organizations during a particular moment: as they were actively designing their response to generative AI. This timing matters. When organizations establish new governance structures and mechanisms, choices that might otherwise remain implicit become visible. The approval-first versus enablement-first distinction exists in many organizational contexts, but observing organizations as they construct AI governance makes visible the default choices they are making, often unconsciously, and the immediate consequences that follow. The patterns described reflect recurring dynamics observed across multiple organizations during this governance formation process.

Two Defaults, Two Paths

Approval-first organizations treat permission seeking as the default. The governing question is: "Do we allow this AI use to proceed?" The burden of proof falls on demonstrating safety before action. When boundaries are unclear, the default response is to require approval.

Enablement-first organizations treat safe-by-design as the default. The governing question is: "How do we design conditions so this type of AI use can proceed safely?" The burden falls on designing conditions that enable safe action. When boundaries are unclear, the default response is to clarify what is permitted and provide infrastructure that makes safe use possible.

Both approaches aim to manage risk and enable innovation. But the principle, approval-first versus enablement-first, produces strikingly different dynamics.

What I Am Observing In Approval-First Organizations

A global professional services firm illustrates the approval-first pattern. After establishing a central AI governance office, leaders announced that the use of AI tools required approval. People throughout the organization had been experimenting with ChatGPT to draft documents, interpret data, and handle routine work. Once approval became the default, the informal experimentation that had been happening broadly across the organization largely disappeared from view.

Some formalized their experiments into proposals that entered approval queues. They found that routine work now required review processes that could take weeks. Many others continued using consumer AI tools, but this work remained invisible to the governance office meant to coordinate it. The structure created to enable AI adoption had instead fragmented both the work and the learning.

I've observed similar dynamics across organizations defaulting to approval-first thinking:

Work tends to become request-driven rather than initiative-driven. People route ideas upward rather than moving forward with experiments. Progress slows to the pace of review cycles. Experimentation does not stop, but its visibility and coordination become constrained.

Learning tends to stay isolated rather than circulating. People continue exploring within their domains, but discoveries remain local. In one financial services firm, a senior AI leader described how analysts in three separate departments had each independently developed effective approaches for using AI in contract analysis. None knew others were developing similar capabilities. With the senior AI leader focused on reviewing major proposals, these smaller experiments fell outside the governance process entirely and never surfaced for coordination. Valuable learning remained trapped.

Strategy often disconnects from practice. Leaders feel pressure to define compelling AI direction for boards and investors. But when experimentation happens underground or moves slowly through approval channels, leadership lacks the ground truth needed to

inform strategy and therefore struggles to shape strategy that reflects how AI is actually creating value in practice. Plans diverge from the messy reality of actual use. The gap between what leadership thinks is happening and what is actually happening widens.

The approval-first default, however well intentioned, tends to make organizations learn more slowly than AI evolves. When rapid technology evolution demands fast learning cycles, approval mechanisms slow precisely when speed matters most.

What I Am Observing In Enablement-First Organizations

Some organizations are taking a different approach. Rather than focusing on approving AI use requests, they double down on accelerating organizational learning at scale while managing risk. They default to enablement-first thinking: designing conditions that make safe, productive experimentation possible without requiring case-by-case permission.

Organizations using enablement-first governance concentrate on three domains. All three embody safe-by-design thinking. They build governance into infrastructure and conventions rather than merely verifying it through approval processes.

They build shared foundations. These include secure places to experiment, clear conventions that specify what is permitted, and common building blocks that reduce friction for people starting work. Clear conventions eliminate the uncertainty that makes people hesitate or seek permission unnecessarily. When people know what is allowed and have safe infrastructure available, they can proceed with confidence.

They create routines that help learning circulate. These provide regular opportunities to compare insights, share findings, and reflect on signals. These aren't overly formal knowledge management systems. They are lightweight practices that make work visible and help discoveries move across organizational boundaries. The routines create opportunities to share learning and surface valuable work regardless of whether it began through formal channels.

They help turn local insights into repeatable practices. When an individual or team finds something effective, governance helps make it explicit so others can adopt it or build on it without losing what made it work. This isn't standardization that removes judgment. It is scaffolding that helps successful practices spread while adapting to local contexts.

This enablement-first approach addresses the challenges that emerge under approval-first thinking. When boundaries are unclear, people don't hesitate because they have clear conventions. When approval processes slow work, shared foundations allow many forms of experimentation to proceed safely without requiring prior permission. Even when experimentation begins informally, shared infrastructure and learning routines increase the likelihood that it becomes visible and can be coordinated, recognized, and scaled.

This points to an important difference in what each governance default centralizes and what it distributes. Enablement-first governance centralizes the design of conditions for safe experimentation — infrastructure, conventions, boundaries — while distributing the experimentation itself broadly across the organization. Approval-first governance centralizes the decision about who may experiment and on what, but because accessible AI tools make experimentation easy regardless of formal approval, much of the actual experimentation proceeds without coordination. The practical consequence is that enablement-first governance makes experimentation more visible, because it occurs within designed conditions. Approval-first governance often makes experimentation invisible by default, because it occurs despite, rather than through, the governance structure.

Enablement-First In Practice: Walmart's Approach

Walmart's response to generative AI exhibits the characteristics of enablement-first governance in practice. As the company expanded internal use of generative AI, leadership faced the foundational choice: default to approval-first (requiring permission before pursuing AI work) or default to enablement-first (designing conditions that enable employees to experiment safely within clear boundaries while controlling risk).

The company chose the enablement-first path. Rather than creating approval processes to decide which ideas could proceed, leadership invested in shaping conditions under which safe experimentation could begin by default.

Walmart built on Element, its machine learning platform, to provide shared infrastructure rather than create an approval layer.⁵ Element gave data scientists, engineers, and teams direct access to reusable AI models and capabilities with governance standards and safety controls embedded from the start.^{6, 7} The infrastructure itself embodied the governance by design.⁸

The company created GenAI Playground as a protected space where employees could experiment directly within clear boundaries.^{9, 10} Rather than asking permission for each experiment, people could try approaches knowing the environment was designed for safe exploration. The boundaries were clear, and within them, action didn't require approval.

Walmart introduced listening sessions to gather feedback from employees experimenting with AI tools.^{11, 12} This stimulated the learning cycle that approval-first governance often lacks. Employees shared what they were trying, what they learned, and what seemed worth scaling. Discoveries could travel.

Through this enablement-first approach, direction-setting emerged from practice rather than preceding it. Leadership didn't need to approve every initiative because the infrastructure, conventions, and learning routines were doing the coordinating work. As the platform matured, Element expanded available models, refined governance standards, and adjusted controls — consistent with an approach that adapts governance based on what is learned in day-to-day practice.

The work of AI leadership became connecting experimentation at scale, helping learning travel, and supporting what proved effective, rather than deciding what could proceed.

The Choice Organizations Face

Many organizations establishing AI governance today default to approval-first thinking, often without explicitly choosing it. It feels responsible, appears to manage risk, and seems to provide control.

But early patterns suggest approval-first governance can unintentionally create the very coordination and learning challenges it aims to solve. It fragments experimentation, drives valuable work underground, and slows the learning organizations require to adapt. Meanwhile, enablement-first governance demonstrates that organizations can accelerate learning while managing risk. They do this not by verifying safety through approvals, but by building safety into infrastructure and conventions through safe-by-design principles.

This is not to suggest that approval-first governance is never appropriate. In contexts involving sensitive personal data, regulated activities such as clinical decision support, or AI applications where errors carry irreversible consequences, approval mechanisms serve a necessary function. The argument here is about the default, not the exception. Where organizations apply approval-first thinking as the default across all AI use, they risk suppressing the broader organizational learning that the majority of low-risk experimentation could safely generate.

The choice also matters because governance defaults tend to become self-reinforcing and path dependent.¹³ Once approval-first patterns establish themselves, they tend to deepen: approval processes attract additional compliance requirements, people learn to route ideas through formal channels rather than experimenting, and the organization's capacity for distributed experimentation atrophies. Reversing these patterns requires dismantling not just processes, but the expectations and habits that have formed around them. This is a far heavier lift than choosing the right default from the start, because early governance choices shape the learning routines and expectations that organizations build around AI use.

Questions For Leaders

If your organization is establishing or evolving its AI governance, consider these diagnostic questions:

What is your actual default? When an ambiguous AI use case arises, what happens? Do people seek permission (approval-first) or proceed within understood boundaries (enablement-first)? Your stated policy may differ from your operating default.

Where does the burden of proof fall? Must people demonstrate that AI use is safe before proceeding (approval-first)? Or must governance demonstrate why a particular use requires special approval (enablement-first)?

What is the exception and what is the rule? In your organization, is exploring with AI the exception that requires justification, or is it the rule that requires clear boundaries?

Are you building safe-by-design? Does your governance embed safety into tools, infrastructure, and conventions so that unsafe use becomes difficult by design? Or does it verify safety by reviewing each proposed use?

What is happening outside the line of sight? Where are people already using consumer AI tools or unsanctioned workflows? What would make that work both safer and more visible for organizational learning?

Where could you shift toward enablement-first? In which domains could AI governance be redesigned around clear conventions and learning circulation? And where do genuinely high-risk contexts still require approval mechanisms?

Looking Ahead

These governance choices are being made now, often implicitly, as organizations respond to generative AI's accessibility. They may shape AI adoption trajectories for years. Organizations that default to approval-first thinking may find themselves struggling with

the very challenges their governance was meant to address: slow learning, fragmented innovation, and gaps between strategy and practice.

Organizations that default to enablement-first thinking position themselves to learn at the pace AI advances demand. They demonstrate that governance can enable rather than constrain. They achieve this not through less governance, but through governance designed around different questions.

The choice is being made with each decision, with each process designed, with each question asked when something new arises.

The question isn't whether to govern AI. The question is: What learning dynamics will your governance design create?

References

1. Thomas H. Davenport and Nitin Mittal, "Why Your Company Needs a Chief Data, Analytics, and AI Officer," *Harvard Business Review*, December 2025.
2. Ethan Mollick, *Co-Intelligence: Living and Working with AI* (New York: Portfolio, 2024).
3. World Economic Forum, in collaboration with PwC, *Leveraging Generative AI for Job Augmentation and Workforce Productivity: Scenarios, Case Studies, and a Framework for Action*, Insight Report (Geneva: World Economic Forum, 2024).
4. Walmart Global Tech, "**Walmart's Element: A Machine Learning Platform Like No Other**," March 14, 2024.
5. Adrian Bridgwater, "**Shining a Light on Shadow AI**," *Forbes*, July 7, 2025.
6. Walmart, "**Walmart Unveils New AI-Powered Tools to Empower 1.5 Million Associates**," *Walmart Corporate Newsroom*, June 24, 2025.
7. Louis Columbus, "**Four Big Enterprise Lessons from Walmart's AI Security: Agentic Risks, Identity Reboot, Velocity with Governance, and AI vs. AI Defense**," *VentureBeat*, August 21, 2025.
8. Louis Columbus, "**How Walmart Built an AI Platform That Makes It Beholden to No One (and That 1.5 Million Associates Actually Want to Use)**," *VentureBeat*, June 24, 2025.

9. Kim Souza, “**Walmart Expands Generative AI to Include a ‘Playground’ for Employees,**” *Talk Business & Politics*, June 13, 2023.
 10. Andrew Budd, “**How Walmart Is Empowering Developers with AI,**” *Walmart Global Tech Blog*, October 22, 2024.
 11. Jennifer Moore, “**Walmart AI: Retailer Taps Employees to Find GenAI Use Cases,**” *TechTarget (SearchCIO)*, September 15, 2023.
 12. Walmart, “**Walmart’s Expanding One-of-a-Kind Associate GenAI Tool to 11 Countries in 2024,**” *Walmart Corporate Newsroom*, January 9, 2024.
 13. Jörg Sydow et al., “**Organizational Path Dependence: Opening the Black Box,**” *Academy of Management Review* 34, no. 4 (2009): 689–709.
-



Stijn Viaene [Follow](#)

Stijn Viaene is a professor of information systems and digital business at Vlerick Business School and KU Leuven. He is the academic lead behind the 4V framework of digital transformation leadership, rooted in leadership-as-practice and paradox management theory research on how organisations adapt and learn under digital turbulence. His latest book is *The Digital Leadership Practice Test: Learning to Think and Act Smarter with Digital and AI* (2026).