



© TechNova Graphics/AdobeStock

LIGHTS - Innovation/Entrepreneurship

AI won't just change how your company works
— it may change who owns its intelligence —
What every C-suite leader must understand
about strategic sovereignty in the age of AI

ESCP Impact Paper No.2026-17-ENG

Davide SOLA, Giovanni SCARSO BORIOLI, Roberto QUAGLIA
ESCP Business School

AI won't just change how your company works — it may change who owns its intelligence — What every C-suite leader must understand about strategic sovereignty in the age of AI

Davide Sola*, Giovanni Scarso Borioli*, Roberto Quaglia*

ESCP Business School

Abstract

As companies adopt AI at scale for speed, efficiency, and competitive parity, few are examining what they may be quietly giving away. This paper asks what happens to an organization's institutional intelligence when AI becomes ubiquitous inside it. Conceptually, we argue that the defining risk of widespread AI adoption is not automation or job displacement but the gradual erosion of strategic sovereignty: an organization's control over its own reasoning, assumptions, and decision logic. Because intangible assets now represent some the value of 92% of S&P 500 market capitalization, this erosion carries a balance-sheet dimension that boards have not yet priced. Drawing on primary C-suite research and emerging corporate evidence, we identify three compounding dynamics: the hidden transfer of institutional logic to platform providers; the erosion of organizational commitment as AI mediates decisions; and the loss of strategic coherence in the context of local, uncoordinated deployment. For managers, we prescribe three governing principles: codify an explicit institutional layer of reasoning; bind AI to that layer with humans at the centre of strategic decisions; and develop agents locally to retain sovereignty over the most critical reasoning capabilities. Institutional intelligence is the emerging strategic asset class: the structured, accessible and traceable representation of how an organization thinks.

Keywords: strategic sovereignty, institutional intelligence, decision architecture, AI Governance, organizational learning

*Professor, ESCP Business School

ESCP Impact Papers are in draft form. This paper is circulated for the purposes of comment and discussion only. Hence, it does not preclude simultaneous or subsequent publication elsewhere. ESCP Impact Papers are not refereed. The form and content of papers are the responsibility of individual authors. ESCP Business School does not bear any responsibility for the views expressed in the articles. Copyright for the paper is held by the individual authors.

AI won't just change how your company works — it may change who owns its intelligence — What every C-suite leader must understand about strategic sovereignty in the age of AI

The question everyone should be asking

In the first quarter of 2026, an estimated \$2 trillion in enterprise software market capitalization was erased in a matter of weeks. A new class of AI agents had arrived: capable of executing multi-step cognitive tasks within familiar workspaces such as MS Office and Windows file explorer. Knowledge workers suddenly had an army of co-workers available to collaborate seamlessly and accelerate their work. They did not need to learn a new system; the system adapted to the way they already worked. Investors concluded that if AI could do the work directly, software might become redundant, and analysts dubbed the sell-off the SaaSocalypse (Bloomberg, 2026).

That reaction focused on vendors. Executives in operating companies should be asking a more fundamental question.

What happens to your organization's institutional intelligence when AI becomes ubiquitous inside it?

The answer has a balance-sheet dimension that boards have not yet fully priced. According to Ocean Tomo's 2025 Intangible Asset Market Value Study, intangible assets now constitute approximately 92% of S&P 500 market capitalization. In 1975, that figure was 17%. Corporate worth has migrated from what can be touched to what can be thought (Ocean Tomo / J.S. Held, 2026).

For most knowledge-intensive firms, the gap between book value and market value, recorded as goodwill, represents accumulated institutional reasoning: the tacit knowledge, decision heuristics, and institutional memory that make a firm more than the sum of its tangible assets. When that capacity migrates outward to AI platforms, a meaningful share of that value can migrate with it. This is not merely a philosophical concern; it is, increasingly, a balance-sheet risk.

"When I think about what we're actually worth, beyond the tangible assets on our books, it's the accumulated judgment of this organization. The way we assess risk, the way we make trade-offs, the institutional memory of what has worked and what hasn't. If that lives inside an external AI system, it's no longer ours. I'm not sure our auditors have a line for that yet. But they should."

— CEO, financial services firm

AI adoption is accelerating regardless: 55% of large European companies are already deploying AI technologies, up from 8% of all firms in 2021 (Eurostat, 2025). McKinsey finds that 88% of organizations globally now use AI in at least one business function (McKinsey, 2025). Most leaders have approached this through a lens of competitive necessity. What is missing is a strategic question: whether the way AI is being adopted preserves institutional intelligence, or quietly depletes it.

That distinction is subtle. It is also decisive. The greatest risk of widespread AI adoption is not job loss, and it is not automation. It is the gradual erosion of strategic sovereignty.

Strategic sovereignty means your organization owns and controls its strategic intelligence: that assumptions are explicit and traceable, trade-offs are documented, decision pathways are structured, and institutional memory compounds over time. It means AI operates inside your governance, not around it.

Without strategic sovereignty, individual intelligence (human or artificial) may become abundant but institutional intelligence deteriorates.

This erosion unfolds across three distinct but compounding dynamics, each already visible in organizations operating with AI at scale, today.

The hidden transfer: when productivity becomes a strategic vulnerability

Organizations have always faced a tension between outsourcing for efficiency and retaining control. Outsourcing or offshoring manufacturing erodes process knowledge and exposes OEMs (Original Equipment Manufacturers) to the rise of low-cost competition. Outsourcing technology R&D boosts short term ROI, but weakens the internal capability to innovate. Each decision appears rational in isolation. Only in aggregate does the strategic cost become visible.

AI adoption is replicating this dynamic at a deeper level and at a faster pace. When organizations deploy AI tools across legal, finance, marketing, operations and strategy, they create a structural dynamic few executives are currently tracking: they are gradually training commercial models on how their companies reason.

Every interaction with an external AI platform contributes to pattern formation. Even when individual data are protected, aggregate interactions allow platforms to learn how industries reason: how contracts are structured, how risks are assessed, how strategic options are framed and how choices are made. Over time, external platforms accumulate a cross-industry map of institutional intelligence that no single firm possesses about itself.

“We’ve been so focused on what AI can do for us that we haven’t asked what we’re teaching it about us. Every time our teams use these tools for strategy work, we’re essentially externalizing how we think. That worries me more than any data breach.”

— CEO, European engineering company

The risk is not that confidential files leak. The risk is that the “inner grammar” of your industry becomes external property (Pignataro, 2026). That is not a technical vulnerability. It is a strategic one.

In March 2023, engineers at Samsung Semiconductor pasted proprietary source code, internal meeting transcripts, and test data directly into ChatGPT on at least three separate occasions. Samsung responded by banning generative AI tools on company devices. What Samsung experienced was not only a security incident. It was a sovereignty failure: the structured reasoning and professional judgment that semiconductor engineers accumulate over careers was being externalized without any governing architecture in place.

A second dimension of the sovereignty risk is subtler still, it is the fact that companies can give away a lot in exchange for absolutely nothing. Indeed, while the productivity gains from AI are real, they are equally available to every company in your industry. A tool that can be licensed by every competitor cannot, by definition, be a source of sustainable competitive advantage. Organizations that outsource their reasoning to the same external platforms will, over time, converge on similar strategic conclusions. What AI commoditizes is not just cost or speed. It commoditizes the reasoning patterns themselves. Empirical evidence confirms this: Csaszar, Ketkar and Kim (2024) demonstrate that LLMs tend to reproduce conventional choices, reducing the novelty and uniqueness of strategic ideas across firms. In Barney's (1991) terms, a resource that is equally available to all competitors fails the test for sustained competitive advantage; institutional intelligence is the resource that passes it.

Strategic sovereignty is not only about protecting what you have. It is about preserving your ability to remain different.

For executives: every prompt written into an external platform is a small deposit of your reasoning into a shared industry account. That distinctiveness leaks, one interaction at a time.

The commitment problem: how AI erodes ownership of strategic decisions

The second risk is less visible, and more insidious. As organizations delegate cognitive work to AI, including analysis, synthesis, scenario generation and even draft strategy documents, something changes in the relationship between humans and decisions. People become recipients of outputs rather than reasoning agents.

Commitment to a strategic direction is not a rational product of analytical quality. It is a social product of participation. Decades of research on organizational change confirm that people commit to decisions they helped make: assumptions they were asked to test, trade-offs they were asked to weigh, logic they were invited to challenge. This is the mechanism by which strategy becomes action.

When AI generates the analysis and frames the recommendation, the human role shrinks. Leaders who inherit a strategy they did not reason through are not strategically committed. They are strategically compliant. And compliance does not produce the organizational energy and adaptability required to execute in the face of adversity, volatility, disruption, or resistance to change.

This dynamic will be familiar to anyone who has worked with management consultants. A strategy developed externally can be analytically rigorous, well-evidenced, and entirely correct. Yet organizations have learned, often painfully, that recommendations that arrive from outside the room tend not to survive the first serious obstacle. People do not fight for conclusions they did not reach. They do not defend trade-offs they were not asked to make. The consulting profession has grappled with this for decades, developing increasingly elaborate mechanisms for building client "buy-in" precisely because the ownership problem is structural, not incidental. The truth is, a large share of the knowledge generated by the consulting team within a strategic assignment remains tacit, and disappears when the consulting team leaves, as the strategic recommendations need to be simple to understand and communicate.

AI makes this problem structurally worse, not merely comparable. A consulting engagement is episodic: it has a beginning and an end, and leaders know the recommendation came from outside. AI is continuous and often invisible. It is woven into the daily fabric of analysis, synthesis, and decision support. Leaders may not even register that their reasoning has been progressively mediated, that the assumptions they are working from were shaped by a model, or that the options they are choosing between were pre-framed by an algorithm. The result is not just a lack of ownership. It is a loss of awareness that ownership was ever at stake. If the consulting engagement diluted commitment intermittently, AI risks doing so continuously and at far greater scale.

The danger is not that AI-produced strategy is poor. It may be excellent. The danger is that it remains external to the organization: technically impressive, but not owned operationally. Strategy succeeds only when people commit to it. Commitment is emotional and social, not computational.

For executives: a recommendation no one reasoned through is a recommendation no one will defend. A strategy fully formed by a model invites compliance, not ownership. The more AI reasons for you, the less your organization is committed to it.

The loss of coherence: how AI optimizes locally and fragments strategically

The third risk is already visible in organizations that have moved fastest on AI adoption. When AI tools are deployed function by function, each deployment is locally rational. Marketing refines targeting. Finance revises forecasts. Operations redesigns routing. But without a shared set of integrated strategic assumptions governing how each function's outputs connect to the whole, the functions can rapidly drift apart.

Kenny and Oosthuizen (HBR, September 2025) document a bank where risk management and marketing AI tools reached opposing conclusions about the same customers: risk flagged them as too risky, while marketing targeted them for growth. The problem was not the algorithm. It was the absence of an integrative strategic logic governing its use.

"We had five different functions all using AI to make better decisions. The problem was that they were making better decisions for themselves. No one was asking whether those decisions added up to anything coherent at the company level. We got faster at going in different directions."

— Chief Strategy Officer, global consumer goods company

The evidence suggests this problem is widespread. BCG's 2025 study of more than 1,250 firms found that 60% reported little to no value from AI despite substantial investment, with only 5% achieving value at scale. McKinsey finds that over 80% of organizations report no significant bottom-line EBIT impact at the enterprise level. In both cases, the diagnosis is the same: not because the technology failed, but because the governing architecture required to connect capability to coherent strategy was absent.

How can organizations maximize the benefits of AI without incurring the loss of strategic sovereignty, organizational commitment, and strategic coherence?

For executives: five functions each optimizing locally can pull a company apart faster than the total absence of AI. Local intelligence, ungoverned, becomes collective incoherence.

Three governing principles: what forward-looking organizations are doing

The firms navigating this challenge most effectively are not those with the largest AI budgets. They are those asking a different question: not “How do we use AI?” but “How do we ensure AI operates within our strategic logic rather than replacing it?” Three governing principles characterize their approach.

1. Codify your institutional knowledge.

AI systems are exceptionally good at optimizing for defined objectives. They cannot ensure that the assumptions behind those objectives remain valid and coherent, are shared across the organization, or stay intelligible to leaders who inherit decisions made before their tenure.

Forward-looking organizations are codifying their institutional knowledge: a structured, explicit, continuously evolving representation of their strategic logic. It is what Nonaka and Takeuchi (1995) describe as the externalization of tacit knowledge into organizational knowledge, and what Senge (1990) operationalizes as the discipline of surfacing and testing the mental models that govern organizational behaviour. This includes documented assumptions, traceable trade-offs, and decision records that preserve the reasoning behind major choices, not just the choices themselves. Every significant AI output connects to a human-owned reasoning architecture that can be interrogated, challenged, and updated.

Thomson Reuters’ Westlaw Edge offers an instructive analogue. The platform sits atop a rigorously curated legal knowledge graph of statutes, cases and regulations, with auditable provenance. AI modules reason over this structure, compressing days of research into minutes while preserving the doctrinal logic that makes outputs trustworthy. Business strategy currently lacks an equivalent structure. Building one is the defining strategic infrastructure challenge of the next decade.

For executives: before connecting a single AI tool, write down how you take decisions and why.

2. Connect AI models to your institutional intelligence, and place humans at the centre of strategic decisions and choices.

Most organizations manage AI adoption as a portfolio of tool decisions: which platforms to license, which use cases to prioritize. This is necessary but not sufficient.

The institutional layer must be the governing architecture to which all AI models are bound. This means requiring explainability for high-stakes decisions: every significant AI inference must trace back to the institutional logic your organization owns and controls, not to patterns absorbed from external training data. It means structuring AI as a set of

utilities, fast and effective at applying your logic to your data, generating signals, synthesizing context and surfacing options for human decision-makers.

It also means deciding in advance which decisions AI can influence, which require human review, and which remain entirely human. This is a critical governance design task. It is how leaders reclaim authorship and ownership over organizational reasoning, and how strategic sovereignty is defended in practice.

This may require a dedicated function, an Office of Strategic Intelligence or equivalent, with explicit ownership of the institutional intelligence layer and authority to set standards for AI deployment. It is not a technology function. It is a strategy function.

Such governance design is enabled by the previous layer, as it is only when leaders engage with an explicitly codified reasoning logic, surfacing and testing assumptions, challenging trade-offs, refining logic, that they can commit to choices, and become authors of strategic choices rather than recipients.

For executives: decide in advance which decisions AI may influence, which need your review, and which remain entirely yours.

"We've changed a number of key executives. Each time, we lose the strategic context: the reasoning behind key decisions, the assumptions that were tested and rejected, the trade-offs that were deliberately made. If AI is running more of our strategy processes, and that institutional knowledge lives in the model rather than in our organization, the next transition will be even more disorienting."

— **Former CEO, large FTSE managed services company**

3. Develop agents locally, fine-tuned on your logic and your data. Keep your most important asset private.

The most strategically consequential principle: keep the most critical elements of your institutional reasoning architecture out of commercial AI platforms entirely. Develop AI agents and fine-tuned models locally, leveraging open-source foundation models trained on your own data and your own institutional logic.

The distinction matters architecturally. A commercially-hosted AI tool can consume the outputs of your institutional knowledge graph without ingesting the graph itself. The AI reasons with the signals your layer produces, including insights, options and risk flags, without learning the logic that generated them. That boundary is a strategic asset.

The practical barriers to doing this have fallen sharply. Platforms such as Hugging Face now host tens of thousands of open-source foundation models, many of them competitive with commercial alternatives on specialized tasks, available at no licensing cost and fully deployable within an organization's own infrastructure. Models such as Meta's Llama family, Mistral, and a growing range of domain-specific variants can be fine-tuned on proprietary data without a single token of that data leaving the organization's environment. What once required the resources of a large technology company is now accessible to any organization with a modest data science team. The cost of local model development is falling rapidly. The cost of sovereignty failure, namely strategic convergence with competitors, erosion of distinctive reasoning capacity and

organizational commitment that withers because leaders no longer own the reasoning, is not.

For executives: let AI reason with the outputs of your logic, never with the logic itself.

The principles operate at any scale – an illustrative case

Consider a successful family-owned pizzeria that, after thirty years of operation, decides to use AI to help manage menu pricing, ingredient ordering, and customer relationships. The scale is modest. The governance challenge is identical to that of a multinational.

Before connecting any AI tool, the owner does something that many large organizations have not yet done: she writes down what makes her business different. Why the dough recipe is what it is. Why certain ingredients are non-negotiable even when cheaper substitutes exist. Why some menu items are kept despite low margins, because they anchor customer loyalty and define the identity of the place. These are decisions she has made over thirty years, mostly in her head. Now they are explicit. In effect, she has built a small knowledge graph of her strategic logic, and she places it in a protected environment she controls entirely. The AI tools she is about to deploy will make recommendations; she wants those recommendations evaluated against her logic, not the aggregate logic of every pizzeria the model has ever seen. Her accumulated reasoning never passes through the commercial model. The AI reasons with the outputs of her institutional layer; it never ingests the layer itself.

She then makes a second set of decisions that most organizations skip. She draws a clear line between the choices AI can influence, the choices that require her review, and the choices that remain entirely hers. Ingredient ordering and demand forecasting: AI can handle those. Menu pricing and supplier changes: she wants to see the recommendation before anything moves. Which dishes define the soul of the restaurant: that question never goes near an algorithm. When her son suggests a new AI scheduling tool, her first question is not whether it works well. It is which decisions it is allowed to touch. This is governance at the level of decisions rather than tools, and it is how strategic sovereignty is protected in practice.

When her son eventually takes over, he inherits not only the ovens and the recipes but the documented reasoning: why decisions were made, what was tried and rejected, what trade-offs were preserved deliberately. The institutional layer does not disappear with the founder; it compounds. When he opens a second location, then a third, the codified logic of the format is the asset that allows him to scale deliberately rather than reconstruct painfully. He is not starting from scratch. He is building on sovereign ground.

The scale of the institutional layer changes. The principles that protect it do not. Sovereignty is not a defensive posture. It is the foundation on which deliberate growth becomes possible.

Conclusion

Artificial intelligence is transforming how organizations operate. But the deepest transformation is not technological. It is architectural. The substitution of AI for institutional reasoning, even when gradual and well-intentioned, redistributes strategic control in ways that balance sheets do not yet capture and governance frameworks have not yet addressed.

Three compounding dynamics define the risk. The hidden transfer: organizations externalize their institutional grammar to platforms that serve entire industries, surrendering the reasoning-based distinctiveness that goodwill represents. The commitment problem: as AI mediates the reasoning behind decisions, leaders become compliant rather than committed, and compliance does not survive adversity. The coherence problem: function-by-function optimization, without a governing architecture, makes organizations faster at going in different directions.

Strategic sovereignty is the answer to all three. It means owning your strategic logic: making assumptions explicit and traceable, documenting trade-offs, structuring decision pathways, and ensuring that institutional memory compounds over time inside the organization, not inside a vendor's model. The three governing principles are the practical means of building and protecting it.

The evidence is unambiguous. 55% of large EU enterprises are already deploying AI (Eurostat, 2025). Yet 60% report little to no value from that deployment (BCG, 2025), and over 80% report no significant enterprise-level financial impact (McKinsey, 2025). The gap is not technological. It is architectural.

The dominant misconception is that competitive advantage will belong to the firms with the most advanced models. Models will be commoditized. The real advantage will belong to firms that define their decision logic clearly, institutionalize strategic reasoning, maintain sovereignty over their intellectual architecture, and integrate AI without surrendering control.

In an age of abundant intelligence, the scarcest resource is not computation. It is coherence. And coherence must be built, not bought.

The question for every executive reading this is not whether to adopt AI. That decision is already made by competitive necessity. The question is whether your organization is building the institutional architecture that ensures AI amplifies your strategic intelligence rather than replacing it.

The firms that answer that question successfully will not just survive the AI transition. They will define it.

References

- Barney, J. (1991). Firm Resources and Sustained Competitive Advantage. *Journal of Management*, 17(1), 99–120.
- BCG (2025). The Widening AI Value Gap. BCG Build for the Future Global Study (n = 1,250). Boston Consulting Group, September 2025.
- Bloomberg (2023). Samsung Bans ChatGPT Use by Staff After Data Leak. May 2, 2023.
- Bloomberg (2026). What's Behind the 'SaaSocalypse' Plunge in Software Stocks. *Bloomberg Technology*, February 4, 2026.
- Csaszar, F.A., Ketkar, H. & Kim, H. (2024). Artificial Intelligence and Strategic Decision-Making: Evidence from Entrepreneurs and Investors. *Strategy Science*, 10(1).
- Eurostat (2025). Use of Artificial Intelligence in Enterprises: Statistics Explained. European Commission, December 2025.
- Kenny, G. & Oosthuizen, K. (2025). Don't Let AI Reinforce Organizational Silos. *Harvard Business Review*, September 18, 2025.
- McKinsey & Company (2025). The State of AI: Global Survey. McKinsey QuantumBlack.
- Nonaka, I. & Takeuchi, H. (1995). "The Knowledge-Creating Company". Oxford University Press.
- Ocean Tomo / J.S. Held (2026). 2025 Intangible Asset Market Value Study. Released February 12, 2026.
- Pignataro, A. (2026). The Wrong Apocalypse. ION Analytics / Mergermarket, February 17, 2026.
- Scarso Borioli, G., Sola, D., Quaglia, R., & Jordan, A. (2025). Towards a Machine-Readable Standard Ontology for Business Strategy. ESCP Impact Paper No.2025-X-X.
- Senge, P.M. (1990). "The Fifth Discipline: The Art and Practice of the Learning Organization". Doubleday.
- TechCrunch (2023). Samsung Bans Use of Generative AI Tools Like ChatGPT After Internal Data Leak. May 2, 2023.