

How boards should govern artificial intelligence

Strong governance accelerates AI deployment rather than slows it down. The STAR framework makes oversight practical and quarterly.



October 5, 2026

By Henk S. de Jong and [Sampsä Samila](#)

While AI is now central to corporate operations, most boards lack a structured way to oversee how their organizations deploy it, manage its risks or capture its value. That gap has gone from being a compliance problem to a leadership problem, which is getting more expensive by the quarter.

The numbers are consistent across surveys from 2024 to 2026: AI adoption is widespread, but most boards have no formal AI oversight structure, and [a majority of directors struggle to keep pace with the technology](#) they are meant to govern. In jurisdictions from the U.S. to Europe and Asia, legal doctrine is clear: Where a risk is central to operations, boards have a fiduciary obligation to put oversight systems in place. For AI, that obligation is already active.

Two ways boards fail in AI governance

Our [paper](#) on how boards should respond to AI, written with Robert Maciejko and Christoph Wollersheim, identifies two recurring failure modes:

- **Under-engagement.** Boards don't understand enough to challenge management, so investment decisions get approved without scrutiny or are deferred indefinitely. The result is value leakage: scattered experiments, no scaling discipline, and competitors, especially AI-native entrants, capture what incumbents leave on the table. This failure is invisible in the short term but devastating over a three-to-five-year horizon.
- **Over-acceleration.** Boards push for rapid deployment before the controls, data infrastructure and operating model are ready. The result is value destruction: algorithmic bias, data breaches, misleading public claims and regulatory action. All this is costly, visible and hard to reverse.

Most organizations are exposed to both problems at the same time. They move too slowly to capture value and too carelessly to manage risk. This is the uncomfortable truth our paper addresses.

5 classic governance responsibilities in an AI context

We propose that boards ground their approach to AI in five familiar governance responsibilities, applied to a new context:

1. Purpose, ethics and compliance

When companies deploy AI to make or shape decisions about customers, employees, risks or resource allocation, they are overseeing choices about what the firm optimizes for and what it's willing to let a machine decide. Corporate purpose must be the starting point for AI governance, not an afterthought.

Compliance is about what is legally permitted. *Ethics* is about what is responsible, given the corporation's values and commitments. *Purpose* is about what is worth doing: which uses of AI serve the corporation's objectives and which undermine them, even if legal and arguably ethical. GenAI may reduce the cost of execution, but the difficulty of deciding what the firm should and should not do remains a board responsibility.

2. Business model and strategy

The strategic question for boards is whether the organization is capturing value commensurate with its investment and risk exposure. For [most companies, the honest answer is not yet](#). The gap between deploying AI and making money from it is real, and boards should be clear about why. Boards should also distinguish value creation from value capture. AI may generate productivity gains, but those gains don't automatically flow to the firm. They may leak to customers, vendors or even competitors.

This leads to the broader strategic question that many boards haven't confronted yet. AI may change what your company does internally, but it can also reshape the competitive environment you operate in: lowering barriers to entry, commoditizing capabilities that incumbents have treated as durable advantages, and compressing the competitive cycle from decades to months.

3. Assets, capabilities and capital allocation

A sound AI strategy can still fail if the organization lacks the assets, capabilities and capital discipline to execute it. In AI, the binding constraints are often organizational, and that's why we treat execution capacity as a separate board responsibility.

GenAI often creates an [illusion of simplicity at the point of use](#). What appears to employees or customers as a straightforward conversational interface may depend on substantial hidden investments in data architecture, evaluation systems, compliance infrastructure and specialist expertise. For boards, this means oversight must extend well beyond adoption

metrics to the assets and governance systems that make reliable deployment possible.

4. Risk profile

AI changes the firm's overall risk profile, across strategy, operations, compliance, reputation, cyber and conduct. The board's job is to understand how AI alters the risks it already governs.

Good AI governance assesses the consequences of each use case, applies controls sized to the risk and maintains an auditable record of the decisions made. Boards should require management to classify AI use cases by risk tier, matching controls, human oversight and audit mechanisms to each tier. High-risk systems need explicit human decision checkpoints, escalation triggers and audit trails designed in before deployment.

5. Leadership selection, evaluation and succession

AI changes what boards should look for in senior executives. Boards should now:

- evaluate whether executives can redesign workflows and operating models, rather than just sponsor pilots.
- orchestrate across business, technology, legal, HR and risk functions.
- distinguish genuine strategic value from hype.
- manage vendor dependency and technical uncertainty.
- communicate credibly with employees about automation, reskilling and the changing nature of work.

The [CEO and top team must lead organizations](#) in which judgment, data, automation and human work are being recombined at speed. Succession planning must reflect these changed requirements.

STAR tool for practical board oversight

With those responsibilities in mind, we introduce the **STAR framework** to make oversight practical and quarterly. Here are the four dimensions of STAR, the questions every board should ask itself, and what boards should be aiming for:

Shareholder value thesis

Board question: Does every major AI initiative have a clear payoff thesis, a named business owner and defined exit criteria?

What good looks like: Management presents AI investments with the same rigor as any capital allocation decision: expected returns, milestones and stopping rules.

Threat parity

Board question: Are our defenses evolving as fast as AI-powered threats, and do we have independent assurance that controls work?

What good looks like: AI creates new security risks: convincing deepfakes, automated vulnerability exploitation, prompt injection. Defenses are tested regularly and independently verified.

Ability

Board question: Do we have the data quality, process standardization, architecture and talent to move beyond pilots?

What good looks like: Real adoption in real workflows, not just licenses purchased. Process redesign underway for top use cases. Talent pipeline filling critical AI roles.

Risk budget

Board question: Have we explicitly defined where AI risk is acceptable, where it's not, and who is accountable when something goes wrong?

What good looks like: Explicit green/yellow/red risk lanes. No-go zones defined. Every high-impact AI system has a named executive accountable for outcomes.

The key is to use STAR as a review discipline, rather than a compliance checklist, to stop failure modes from going undetected quarter to quarter.

Checking for fairness, accuracy and transparency

Governance framed as a constraint is governance that slows organizations down. Effective AI governance works like power steering: It gives you directional control at speed. Organizations that have pre-cleared their AI systems for fairness, accuracy and transparency deploy faster, not slower, because the questions that stall every rollout have already been answered.

Boards that govern AI well are boards that ask the right questions, insist on evidence and hold management accountable for results. AI does not diminish the role of human leaders; it demands more from them than ever.

ALSO OF INTEREST:

[More than half of boards of directors have yet to see significant business value from AI](#)

[Airbus Chair Amparo Moraleda on the governance challenges of AI](#)

[GenAI: Easy to use, harder to manage](#)

[How to put AI to work in your organization](#)

A shorter version of this article was published in the [Center for Corporate Governance](#) newsletter. To sign up to the newsletter, click [here](#).



Sampsa Samila

Professor of Strategic Management and Director of the [AI and the Future of Management Initiative](#). He focuses on how technologies like AI enable new business models, the capabilities needed to deliver value, and how executive mindsets need to change in light of new developments.

www.iese.edu/insight