First 90 Days on an AI-Heavy Board

Idea In Short

Nonexecutive directors joining an AI-heavy board should start with a simple aim: understand where artificial intelligence shapes decisions, communication and risk before they approve more investment or accept management assurances. The first 90 days should focus on mapping exposure, testing governance and checking whether the board receives evidence or only optimistic narratives

Start With Exposure, Not Hype

A new nonexecutive director should begin by asking where AI already affects the enterprise. Many boards spend too much time on pilot programs and future opportunities while missing current exposure in customer service, internal decision support, fraud monitoring, communication, hiring and product features. Board guidance on AI increasingly stresses that oversight must cover the full range of company AI activity, not only flagship initiatives.1

The first useful request is a map. Management should show which business units use AI, which systems are experimental, which are scaled and which involve external vendors. The director should also ask which uses affect regulated decisions, customer outcomes, or public communication. That map reveals whether the board is dealing with contained tools or with a companywide operating layer.

This stage is about exposure before enthusiasm. A director who understands where AI already sits inside workflows can ask better questions about accountability and control. Without that view, the board risks discussing AI as a concept while missing the places where errors, bias, misinformation, or uncontrolled automation can cause damage.

Ask Who Owns What

The second priority is ownership. New directors should ask who in management owns AI governance, model risk, data quality, communication policy and incident response. Many organizations diffuse responsibility across technology, legal, risk, compliance and business teams. That diffusion often creates blind spots where important issues fall between functions.2

A practical board question is: who can stop an AI deployment that creates unacceptable risk? Another is: who signs off on AI-enabled communication that reaches investors, employees, or customers? These questions are simple, but they reveal whether the control environment is real. If answers are vague or split across too many leaders, the board may be facing coordination risk rather than only technology risk.

The director should also test committee alignment. Audit, risk, nomination and governance and technology committees often all touch AI. The first 90 days are the right time to ask whether committee charters clearly assign responsibilities or whether each committee assumes another group is covering the topic.

Test the Quality of Board Information

Boards often receive polished AI updates that emphasize opportunity, investment and momentum. A new director should examine whether the board also receives evidence on incidents, near misses, false positives, overrides, model drift and communication failures. Governance specialists have warned that AI cannot remain a delegated management issue because directors need enough engagement to challenge assumptions and review decision quality.3

One strong early question is whether management reports leading indicators or only headline outcomes. For example, the board should know not only that an AI-enabled process improved speed but also whether error rates rose, whether customer complaints changed and how often humans overrode the system. This helps the board distinguish operational gain from hidden fragility.

The director should also ask to see one full decision trail. Choose a real AI use case and follow it from design through approval, deployment, monitoring and escalation. That exercise often reveals more than ten polished presentations because it shows how governance works in practice rather than in policy language.

Focus on Communication Risk

In AI-heavy companies, communication risk deserves specific attention. AI now shapes drafts for investor messages, employee updates, customer service responses and marketing copy. New directors should ask where AI assists communication and what review standards exist before messages go live. This matters because communication failures can damage trust faster than operational mistakes, especially when errors appear public and avoidable.

The board should ask whether the company maintains one factual narrative about AI use and then tailors that narrative for investors, employees and customers. That approach helps preserve consistency while respecting different stakeholder concerns. It also reduces the risk that separate functions use AI tools to create conflicting messages about the same issue.

A director should further test how the company handles mistakes. If an AI-assisted customer message contains a false claim or an employee note uses the wrong tone during a sensitive change program, who detects it and how quickly does escalation occur? These questions push management to show whether communication controls are substantive or cosmetic.

Build Personal Fluency Fast

New directors do not need to become technical experts, but they do need working fluency. In the first 90 days, they should schedule short briefings on model types, vendor dependencies, data governance and common failure modes. They should also ask management to explain the firm's AI stack in plain language. If explanations remain vague or overloaded with jargon, board oversight will stay weak.

Personal fluency also comes from pattern recognition. Directors should compare AI oversight with earlier waves of digital, cyber and data governance. This helps them see what is familiar, such as third-party dependence or control design and what is distinct, such as probabilistic outputs and content generation. The point is not to import old frameworks unchanged but to avoid treating AI as completely unprecedented.

Finally, the director should decide what good board oversight looks like by day 90. A reasonable goal is not perfect certainty. It is a practical baseline: clear ownership, mapped exposure, evidence-based reporting, escalation routes and a board that can discuss AI in business terms rather than in slogans alone.

In conclusion, a first-time director on an AI-heavy board should spend the first 90 days mapping exposure, clarifying ownership, testing reporting quality and examining communication controls. Those questions build board fluency quickly and help convert AI oversight from broad curiosity into disciplined governance.

- 1AI and the Role of the Board of Directors
- 2AI Governance: A Guide for Boards, Risk and Audit Leaders
- 3Are corporate boards ready to govern AI at speed?

Summary

Directors who ask precise questions early can move AI oversight from slogans to substance. In the first 90 days, the goal is not technical mastery. It is informed skepticism, governance clarity and better board judgment