Nemko Digital

The Five Key Questions to Ask in the Board on AI

When artificial intelligence enters the boardroom, discussions often drift toward technology, tools, or isolated risks.

That is rarely where the real governance questions sit. The most effective board conversations on AI are not technical at all. They are diagnostic.

There are a small number of questions that, when asked well, immediately reveal how mature an organisation's AI governance really is. These are the questions that cut through ambition, PowerPoint assurances, and fragmented ownership. They work in any boardroom setting, regardless of sector or level of AI maturity, because they expose whether AI is being governed deliberately or merely tolerated.

AI is now embedded in core business processes, from customer interaction and pricing to fraud detection, risk assessment, and operational decision making. As a result, AI governance has become a matter of strategic oversight, accountability, and long-term value creation. Boards often ask what they should be discussing when AI appears on the agenda. Experience shows that effective oversight does not require deep technical detail. It requires asking the right questions and insisting that the answers are clear, owned, and defensible.

These five key questions consistently distinguish organisations with mature AI governance from those relying on assumptions. Learn them, use them, and lead the discussion.

1. Who is formally accountable for AI-driven decisions

It is tempting to say that accountability for AI simply sits with "the business." In practice, that answer is too simplistic to be meaningful. The functioning of AI is shaped by a combination of development choices, data dependencies, risk assessments, controls, and operational use. Each of these elements often involves different roles, teams, and decision makers.

Mature organisations recognise that accountability for AI is therefore layered and deliberate. They are clear on who owns model development, who owns risk acceptance, who owns operational deployment, and who is accountable for outcomes when something goes wrong. First, this way of thinking must be accepted. Making accountability transparent is the next step. Making it scalable and repeatable across use cases is the third.

When accountability remains implicit or spread across committees, escalation paths blur and progress comes to a halt. Governance tends to fail precisely when scrutiny increases, whether from regulators, customers, or the market.

2. Where AI is used today, and where it should be introduced

You cannot govern what you cannot see. One of the most common gaps boards encounter is incomplete visibility into how AI is actually being used across the organisation. Shadow AI, local experimentation, and embedded vendor functionality often mean that AI adoption runs ahead of formal oversight.

Effective governance starts with visibility, but it cannot stop there. Boards need insight not only into current AI deployments, but also into planned developments and expected scaling.

This question connects directly to value creation. Where should AI be used, and where does it genuinely add value? It forces alignment between AI investment, business strategy, and risk appetite, rather than allowing adoption to be driven by convenience, vendor capability, or local enthusiasm. Organisations that answer this well are making conscious choices. Those that do not are letting direction emerge by default.



3. What AI may automate and what must always involve human judgment

This is often framed as an ethical debate, but it is just as much a risk management and control question. Deciding which decisions can be automated and which require human judgment is a strategic boundary-setting exercise.

Boards play a critical role in defining this compass. Without explicit guidance from the top, boundaries tend to be set implicitly through system defaults, efficiency pressures, or technical feasibility. That rarely reflects the organisation's true risk tolerance or values.

Clear decision boundaries help protect individuals, safeguard the organisation, and provide operational clarity to teams designing and deploying AI systems. Where these boundaries are vague, risk accumulates quietly and responsibility becomes harder to assign. The OECD AI Principles provide a foundational framework for establishing these ethical and governance boundaries.

4. How do we stay in control of AI over time and prove it when challenged

AI governance does not end at deployment. Models evolve, data shifts, and usage patterns change over time. Less mature organisations treat approval as the end point. More mature ones treat it as the starting point of continuous oversight.

Boards should ask how the organisation stays in control at scale. That includes how performance is monitored, how deviations are detected, how incidents are escalated, and how decisions are reviewed over time. The real differentiator is evidentiary readiness. When challenged by regulators, customers, or partners, can the organisation demonstrate with evidence how and why an AI system behaved as it did?

At scale, this cannot be done manually. Tooling, structured processes, and clear ownership are what separate confidence from wishful thinking. International standards such as ISO/IEC 42001 provide guidance on establishing AI management systems that enable continuous oversight.

5. How AI risks introduced by vendors and suppliers are managed

This is one of the most frequently overlooked areas of AI governance, despite being one of the most consequential. For most organisations, third parties are already the primary source of AI functionality, and that dependency will only increase.

Outsourcing technology does not outsource accountability. Boards should understand how AI-related risks introduced through suppliers are identified, assessed, and integrated into existing procurement, risk, and compliance processes. Reliance on vendor assurances or contractual language alone is increasingly insufficient in a regulatory and liability-driven environment.

If this topic is not explicitly on the board agenda, AI risk is almost certainly entering the organisation through the supply chain unchecked. Establishing robust AI management systems helps organisations maintain control over both internal and vendor-supplied AI capabilities.



Why these questions provide the right practical boardroom lens

Individually, each of these questions addresses a distinct dimension of AI governance. Taken together, they form a practical and revealing test of organisational maturity. Boards that can answer them clearly tend to demonstrate deliberate accountability, visibility and intent across AI use, explicit decision boundaries, continuous oversight, and disciplined supplier governance.

Boards that struggle to answer these questions often do not lack ambition or technical capability. What they lack is clarity. In many cases, gaps only become visible after incidents, regulatory scrutiny, or market pressure force the issue. By then, the conversation has already shifted from governance to damage control.

AI governance does not require boards to become technical experts. It requires them to set direction, insist on ownership, and demand evidence. These five questions provide a shared language to do exactly that. The World Economic Forum's AI Governance Alliance brings together global perspectives on how boards and leadership can effectively oversee AI development and deployment.

Bringing them into the boardroom shifts the discussion from reassurance to reality, from intent to execution. If they can be answered consistently and confidently, AI governance is likely robust. If not, they offer a focused and strategic starting point, before governance is tested externally rather than internally.



Visual Summary of the 5 Key Questions to Ask in the Board on AI

About the Author



Bas Overtoom

Bas Overtoom is the Global Business Development Director where he leads global efforts to promote responsible AI adoption at Nemko Digital, working with organizations to operationalize trust, transparency, and compliance in their AI systems. With a strong background in business-IT transformation and AI governance, he brings a pragmatic approach to building AI readiness across sectors.

Nemko Digital