FENWICK

Fenwick Insight

May 27, 2026

AI in the Boardroom: What Directors Need to Know Now

By Marie C. Bafus, David A. Bell, and Wendy Grasso

What You Need to Know:

- Directors have fiduciary duties to implement and monitor AI governance systems, particularly when AI represents a mission-critical business risk. Under the *Caremark* doctrine, sustained or systematic oversight failures may expose directors to personal liability.
- Boards should consider adopting an AI governance framework aligned to business goals, covering (1) strategy and investment, (2) risk assessment and oversight, (3) accountability for AI systems and decisions, (4) stakeholder transparency, and (5) legal and regulatory compliance.
- AI tools have the potential to rebalance the information dynamic between boards and management and enable more independent board oversight, but they require guardrails around confidentiality and privilege.
- To mitigate risks surrounding data retention and discovery exposure from AI tools, boards should consider establishing clear policies governing which AI tools may be used by board members, what data may be shared, how long information is retained, and what security standards the provider must meet.

As artificial intelligence becomes more central to a company's operations and long-term strategic positioning, directors must develop and maintain a robust understanding of how AI impacts their company, including the full spectrum of opportunities, risks, and ethical considerations, and apply the same oversight approach they apply to other areas of board responsibility, including corporate strategy, management performance, regulatory and legal compliance, risk management, and public disclosure obligations. Similar to the framework established for cybersecurity oversight, directors have fiduciary duties to implement and monitor AI governance systems, particularly when AI represents a mission-critical risk to the business. In

fulfilling this duty, boards should work to ensure that adequate governance frameworks, policies, and internal controls are in place to guide the responsible development, procurement, and deployment of AI technologies across the organization.

This article addresses certain key dimensions of board AI governance: director fiduciary duties related to AI, the five pillars of an AI governance framework, responsible AI usage in the boardroom, practical considerations for AI-assisted board operations, and additional governance considerations.

Understanding Board Fiduciary Duties Related to AI

Directors have a duty of care that requires them to act with the diligence a reasonably prudent person would apply when making significant decisions. As AI becomes more critical to a company's efficiency and competitiveness, total inattention to its strategic implications could be viewed as a failure to be informed. Boards must oversee capital allocation related to AI investments and understand how those investments align with the company's long-term strategic goals, risk tolerance, and ethical standards. Courts rarely second-guess business strategy absent gross negligence, so while failing to adopt AI altogether could represent a missed business opportunity, AI-related strategic decisions will generally receive business judgment rule protection. On the other hand, ignoring AI-related risks, including legal compliance, privacy/cybersecurity, intellectual property considerations, and workforce and employment impacts, while the technology is clearly relevant to operations could present an oversight issue.

Under the *Caremark* doctrine, directors may face personal liability for a sustained or systematic failure to implement and monitor adequate oversight systems. In *Stone ex rel. AmSouth Bancorporation v. Ritter* (2006), the Delaware Supreme Court articulated two scenarios giving rise to oversight liability: (1) the director utterly failed to implement any reporting or information system or controls, or (2) having implemented such a system, the director consciously failed to monitor or oversee its operations, thus disabling themselves from being informed of risks or problems requiring their attention. Courts have increasingly extended this doctrine to technology-related risks, and AI oversight may follow this trend, particularly where AI, or how AI will impact a company, is viewed as core or mission critical to the company and its business. The Delaware Court of Chancery's decision in *Firemen's Retirement System of St. Louis v. Sorenson* (2021) built on this principle, recognizing that "as the legal and regulatory frameworks governing cybersecurity advance and the risks become manifest, corporate governance must evolve to address them." The court further noted that "the corporate harms presented by non-compliance with cybersecurity safeguards increasingly call upon directors to ensure companies have appropriate oversight systems in place." Directors should ensure their company maintains active monitoring systems for AI compliance and ethics, as liability in this area often stems from sustained inattention rather than mere bad results.

Furthermore, following *Stone v. Ritter*, the duty to act in good faith is understood as a subsidiary element of the duty of loyalty. A conscious disregard of known AI-related risks, particularly where AI is mission critical, could be characterized not just as a duty of care failure but as a lack of good faith.

Establishing a Clear AI Governance Framework

At its core, AI governance is about overseeing how a company spends money on, invests in, and uses AI tools and deals with AI-related risks. This includes making sure such investments deliver results; that the tools are used safely, responsibly, fairly, and in compliance with the law; and that AI-related risks to the company's technology, business model, or operations are considered and appropriate steps are taken to mitigate those risks. At the board level, this means ensuring that AI efforts support the company's financial, operational, and strategic goals and fit within its broader business strategy.

Key components of an AI governance framework may include:

- AI strategy and investment alignment
- Risk assessment and oversight
- Accountability for AI systems and decisions
- Appropriate transparency with stakeholders
- Compliance with applicable laws and regulations

AI Strategy and Investment Alignment

A strong AI strategy should address the amount of money and resources being devoted to AI, how AI fits into the company's overall business plan, how AI is being used by competitors and where the market is heading, and whether the company's AI practices meet ethical standards with an emphasis on transparency, safety, fairness, and accountability. Boards should also ensure that AI systems are regularly tested, validated, and audited to confirm they are working as intended and delivering value for the investment.

Integrating AI into Risk Assessment and Oversight

AI should be folded into a company's overall enterprise risk management (ERM) framework to identify, assess, and prioritize potential risks. This includes cataloging all AI models, systems, and capabilities, mapping data and AI flows, and identifying and quantifying risk. Directors should work with management to develop strategies and controls to mitigate risk, integrate relevant policies and guidelines into day-to-day processes, and establish metrics and KPIs to measure progress toward strategic goals and validate adherence to AI policies and guidelines.

Importantly, AI-related risks are not uniform, and they may vary significantly depending on the specific use case, the characteristics of the AI system involved, and the industry context in which it operates. A customer-facing generative AI application, for example, may present materially different risk considerations, including reputational, regulatory, and liability exposure, than an internal predictive analytics tool used to optimize supply chain logistics. Similarly, AI systems that involve autonomous decision-making, process sensitive personal data, or operate in highly regulated industries such as financial services, healthcare, or critical infrastructure will generally warrant heightened scrutiny and more robust controls. Boards should therefore tailor the company's risk assessment processes to account for these distinctions, rather than applying a one-size-fits-all approach to AI governance.

Regular reporting and ongoing monitoring may help boards stay ahead of emerging risks and make sure that the company's AI governance practices are still working. Just as with other areas of risk, directors who consistently ignore AI governance may open themselves up to personal liability for failing in their oversight responsibilities.

Promoting Accountability and Ethical Oversight

Clear accountability is fundamental to trustworthy AI governance. While the day-to-day work of implementing AI strategy and managing individual AI systems falls to management, boards are responsible for putting the right structures in place. This includes confirming that management has assigned clear ownership over AI strategy execution and each AI system, establishing ongoing monitoring and audit trails to track how AI-driven decisions are being made and what factors influence them, and creating channels for employees and stakeholders to flag issues and challenge AI-driven outcomes.

Companies may also consider establishing an ethical review board (a dedicated internal body, typically composed of cross-functional representatives from legal, compliance, technology, data science, and other relevant functions) to evaluate whether the company's AI systems and practices align with its ethical standards and values. An ethical review board can strengthen oversight by providing an independent check on management's AI decisions, reviewing proposed AI applications before they are developed or deployed, identifying potential risks related to fairness, bias, transparency, and privacy, and recommending safeguards to prevent reputational and regulatory harm. For boards of directors, an ethical review board at the management level may offer an additional layer of assurance that the company's AI activities are being developed and used responsibly.

Ensuring Transparency and Legal Compliance

Transparency is essential to maintaining trust, both within the organization and with external stakeholders such as investors, regulators, and customers. To build and sustain that trust, boards should ensure that management provides clear disclosures about:

- How AI systems work, how they were developed, and what data they rely on
- How AI is being used to make or inform business decisions
- The known limitations and risks of the company's AI applications

Public companies in particular should be careful not to overstate or misrepresent their AI capabilities, a practice sometimes referred to as "AI washing," and should ensure that their risk factor disclosures accurately reflect the material risks associated with their use of AI. Companies that overstate the role of AI in their business, whether in investor presentations, earnings calls, Securities and Exchange Commission (SEC) filings, or product marketing, risk regulatory action, shareholder litigation, and significant reputational harm.

The SEC Investor Advisory Committee has recommended that the SEC require companies to disclose board oversight of AI, and leading institutional investors are also urging greater transparency on AI usage.

Boards must also ensure that their organization complies with applicable AI laws and regulations. Because the regulatory landscape is evolving rapidly across federal, state, and international jurisdictions, directors should stay informed of emerging legislative and enforcement developments and confirm that management has implemented appropriate policies, information systems, and internal controls to keep the company's AI activities consistent with its legal, regulatory, and ethical obligations. In doing so, boards should be mindful that AI-related compliance exposure can arise not only from internally developed technology but also from tools licensed from third parties or acquired through M&A activity.

Use of AI to Comply with Board Oversight Duties

In addition to overseeing AI-related risks, boards and board members are increasingly turning to AI as a mechanism to assist with their oversight duties. This raises novel questions under *Caremark*.

The first prong of *Caremark* requires a corporation to implement in good faith reporting or information systems or controls that are reasonably designed to detect and escalate risks. Where AI replaces traditional information systems and controls, how does a board make a good faith determination that the AI system is reasonably designed? After all, AI involves complex algorithms that are, by design, autonomous and adaptive, constantly learning and adapting based on the data ingested. Does *Caremark* require directors to have a deep technical understanding of how the algorithm works? No. But neither can a board passively rely on the AI system's integrity merely because it is autonomous and data driven. Rather, as with prior non-AI information systems, *Caremark* requires that a board exercise informed judgment regarding what the AI system is designed to do, how it was chosen, what risks it monitors, whether it is appropriately raising risks, and how to evaluate its performance (including as the model adapts and changes).

The second prong of *Caremark* asks whether, having put an information system or controls in place, directors consciously failed to monitor or oversee its operations, thus preventing them from being informed of risks or problems requiring their attention. In other words, whether a board consciously ignored "red flags." With an AI-based information system, the question is not only whether a board consciously disregards a red flag surfaced by the AI information system, but also whether a board consciously disregards evidence that the system itself is not appropriately surfacing and escalating red flags. For example, given the complex algorithms and adaptive nature of AI systems, an AI information system may be prone to model creep and inherent biases that over time inhibit its effectiveness at ferreting out and raising red flags. Although *Caremark* does not require a perfect system that appropriately escalates every risk, it may require directors to inform themselves as to how the system discovers and escalates risks and what the system qualifies as a risk worthy of escalation, and to take action where there is evidence that the system is not surfacing risks appropriately. In addition, a board should consider periodically reviewing and cross-checking the risks escalated by the AI system against risks and concerns raised by external sources (such as auditors and regulators) and traditional internal control mechanisms to ensure the AI system is appropriately detecting and escalating risks.

Leveraging AI Responsibly in the Boardroom

Aside from using AI as a monitoring system under *Caremark*, AI tools also offer board members a practical way to enhance the quality of their oversight and decision-making. A longstanding governance challenge is that directors depend almost entirely on management for the information they use to evaluate company performance and strategy, a dynamic that may limit the board's ability to form fully independent views. AI has the potential to rebalance that dynamic in several ways.

Directors may use AI to prepare more effectively for meetings by independently analyzing board materials, identifying areas where additional information may be needed, and developing more targeted lines of inquiry.

AI may also support independent benchmarking, allowing directors to compare the company's performance, disclosures, and strategic positioning against publicly available peer data and industry trends rather than relying solely on management's framing.

At a more advanced level, AI-powered tools can synthesize operational and financial data from across the enterprise alongside external market data into accessible formats that help directors focus on material changes, while AI-enabled modeling can help boards evaluate strategic outcomes, test underlying assumptions, and identify and assess key risks to inform decisions on capital allocation and risk tolerance.

These capabilities, however, require careful guardrails. Directors with ready access to independent data and analysis may be drawn toward second-guessing operational decisions that properly sit with management, potentially undermining the collaborative dynamic between the board and the executive team. AI-generated outputs also carry inherent reliability concerns, given that results may appear well-reasoned yet reflect incomplete data, outdated assumptions, or systematic bias, making human verification essential before any AI-assisted insight informs board action. There are also significant confidentiality and privilege considerations: Board materials uploaded to AI platforms may be stored, processed, or used in ways the company cannot fully control, creating risks around proprietary information, corporate recordkeeping, and the preservation of attorney-client privilege.

AI should enhance, not replace, the board's deliberative process, and directors should exercise particular caution when considering the use of AI tools in connection with sensitive matters such as executive compensation, strategic transactions, or legal proceedings.

Practical Considerations for AI-Assisted Board Operations

When boards consider using AI for tasks such as notetaking or meeting logistics, they should be aware of several potential pitfalls.

Current AI transcription tools often have difficulty distinguishing among speakers in a multi-person discussion, which may result in statements being attributed to the wrong director. This could create inconsistencies with other corporate records and complicate matters if the minutes are later scrutinized in litigation or an audit. Even setting attribution aside, accents, pacing,

and other style idiosyncrasies often lead to substantive errors in transcription, and AI is not yet capable of capturing the nuance, tone, and broader context of a complex board discussion.

The mere presence of an AI transcription tool may also alter boardroom dynamics. Directors who know their remarks are being recorded and tagged to them individually may self-censor, dampening the candid debate and constructive dissent on which effective governance depends.

More fundamentally, determining what to memorialize and, just as importantly, what to omit from board minutes requires experienced human judgment, particularly where conversations touch on strategic deliberations, conflicts of interest, executive compensation, or pending or threatened litigation.

Boards must also think carefully about data retention and discovery exposure. AI tools typically operate through cloud-based platforms, and once board materials are uploaded to those environments, the organization may have limited control over how the data is stored, who can access it, and whether it might be used to train the provider's models, potentially placing confidential information in the public domain. The security risks are compounded by the possibility of breaches, unauthorized access, or cyberattacks targeting the AI provider's infrastructure.

From a legal standpoint, information retained by an AI tool may be deemed discoverable by adverse parties in litigation, and privileged communications discussed during a board meeting could lose their protected status if they are captured and stored outside of channels the company controls.

To mitigate these risks, boards should consider establishing clear policies governing which AI tools may be used in board settings, what data may be shared with those tools, how long information is retained, and what security standards the provider must meet. Often, as a practical matter, this means only allowing AI tools that have been approved for use inside the company's systems and infrastructure.

Additional Board Considerations

Boards should engage in strategy and scenario planning with management and evaluate whether the company has the right AI experience and skills at the management level.

AI literacy at the board level is also increasingly important, and directors should pursue education through independent learning, immersive experiences, and engagement with AI advisory councils.

Boards should also consider their oversight structure for AI, whether it be at the board level or designated to one or more committees of the board; determine what information will be presented to the board and by whom; and establish the frequency of AI discussions at the board and committee level.

Key Takeaways

AI offers companies significant opportunities to drive efficiency, strengthen competitiveness, and enhance long-term strategic positioning, and boards themselves may leverage AI tools to improve the quality of their oversight and decision-making. But attaining that value requires a governance infrastructure that matches the technology's complexity and pace of change.

Directors are expected to make reasonably informed decisions, and as AI becomes more central to a company's operations and long-term strategic positioning, total inattention to its implications could be viewed as a breach of the board's oversight duties. The *Caremark* doctrine and its extension to emerging technology risks in the cybersecurity context signal that AI governance is a natural next step, particularly where AI is core to a company and its business. Boards that take a proactive approach by establishing clear accountability structures, integrating AI into enterprise risk management, promoting transparency with stakeholders, and staying current on the rapidly evolving regulatory landscape may position their companies to capture AI's strategic benefits while safeguarding the company against its risks.