SANDBERGPHOENIX

Board Oversight of AI

Is Your Board Behind the Curve on Developing a Framework for Emerging Risks?

By **Mark Williams**, **Vicki Westerhaus**, *and* **Graice Ricci**
Sandberg Phoenix & von Gontard • Saint Louis, Missouri

Artificial intelligence is rapidly becoming embedded in not only our daily individual lives, but also core business operations, from underwriting and pricing algorithms to supply chain management, fraud detection, and regulatory compliance. Yet many boards are behind the curve on developing governance structures for managing AI risk.

A few recent developments illustrate why it is critical to act now:

- the FTC's "Operation AI Comply" initiative continues to target companies for deceptive AI product claims across multiple administrations;[1]
- the EEOC settled its first AI-related employment discrimination case against iTutorGroup over age-discriminatory hiring software;[2]
- a federal court has allowed a nationwide collective action to proceed in *Mobley v. Workday* over allegations that an AI resume-screening platform perpetuates discriminatory patterns at scale;[3]
- at the state level, Colorado's Consumer Protections for Artificial Intelligence Act (SB 24-205) imposes affirmative duties on developers and deployers of high-risk AI systems, with additional states following suit;[4] and
- in the M&A context, a CEO's AI prompts were discoverable and a key piece of evidence in a recent $250 million earnout dispute.[5]

[1]See Fed. Trade Comm'n, "Operation AI Comply" (Sept. 2024) (enforcement sweep targeting deceptive AI claims); *see also* FTC actions against *DoNotPay, Inc.* (Jan. 2025), *Evolv Technologies* (Nov. 2024), and *IntelliVision Technologies Corp.*

[2]*EEOC v. iTutorGroup, Inc.*, No. 1:22-cv-02565 (E.D.N.Y. 2023) (settled Aug. 9, 2023; $365,000 consent decree).

[3]*Mobley v. Workday, Inc.*, No. 3:23-cv-00770-RFL (N.D. Cal.) (motion to dismiss denied July 2024; ADEA collective conditionally certified May 2025).

[4]Colo. Rev. Stat. §§ 6-1-1701 to -1707 (SB 24-205, effective June 30, 2026).

[5]See *Fortis Advisors, LLC v. Krafton, Inc.*, C.A. No. 2025-0805-LWW (Del. Ch. March 16, 2026).

Companies that deploy AI in these contexts should strongly consider implementing monitoring protocols and governance procedures commensurate with the risks presented, as described in more detail in the Practical Recommendations section below. Boards that have not yet developed reporting systems, escalation protocols, and board-level oversight for AI may face increasing fiduciary exposure.

AI and Positive-Law Obligations

Fiduciary duty doctrines like *Caremark*[6] and its progeny already impose obligations on boards and officers to monitor material legal risks, and AI is generating exactly the kind of exposure these doctrines contemplate. *Caremark* claims are strongest where the underlying risk involves violations of positive law, not mere business losses.[7] AI is rapidly becoming a source of just such exposure. In safety-critical domains, the analogy to existing case law is most direct: AI-powered medical devices, autonomous vehicle systems, and AI-driven safety equipment that cause physical or economic harm expose companies to strict liability and failure-to-warn claims.

Where a board fails to oversee AI safety testing and incident-response protocols in these contexts, the logic that courts applied to Boeing's failure to oversee its flight-control software is particularly instructive.[8] Beyond safety-critical applications, the analogy requires more care. AI systems used in hiring, lending, content moderation, and similar contexts present different risk profiles, and the case law is still developing. But the exposure is real and growing: the EEOC's *iTutorGroup* settlement and the ongoing *Mobley v. Workday* litigation demonstrate that existing anti-discrimination statutes apply to AI-driven decisions.

The FTC has brought enforcement actions under Section 5 against companies making unsubstantiated AI performance claims. And Colorado's SB 24-205, which takes effect in mid-2026, imposes impact assessment, transparency, and anti-discrimination obligations on deployers of AI systems that make "consequential decisions" in employment, housing, healthcare, and lending—with the Colorado Attorney General empowered to bring enforcement actions for violations.

AI systems also implicate a range of privacy and data protection statutes. The EU's GDPR, state privacy acts, and the FTC Act all impose obligations on companies that collect, process, and use personal information in AI-driven systems.[9] The *Marriott* data breach litigation illustrates how these risks increasingly call upon directors to ensure appropriate oversight systems are in place—a principle that may extend to AI's novel vectors for data misuse, re-identification, and algorithmic profiling, though the translation from cybersecurity to AI

[6] *In re Caremark Int'l Inc. Derivative Litig.*, 698 A.2d 959, 970 (Del. Ch. 1996).

[7] See *Stone v. Ritter*, 911 A.2d 362, 370 (Del. 2006).

[8] See *In re Boeing Co. Derivative Litig.*, No. 2019-0907-MTZ, 2021 WL 4059934 (Del. Ch. Sept. 7, 2021).

[9] See, e.g., European Union General Data Protection Regulation, Council Regulation 2016/679; Cal. Civ. Code §§ 1798.110, 1798.150; 15 U.S.C. § 45 (FTC Act).

governance is not automatic (discussed further below).[10] Beyond privacy, AI creates exposure under consumer protection statutes, intellectual property law, and, in safety-critical domains, product safety regulation. These are the positive-law predicates that make *Caremark* claims viable.

The *Caremark* Duty Applied to AI

Chancellor Allen's 1996 opinion in *In re Caremark International Inc. Derivative Litigation* established that a board may not satisfy its obligation to be reasonably informed about the corporation without assuring itself that information and reporting systems exist that are reasonably designed to provide senior management and the board with timely, accurate information about the corporation's compliance with law and its business performance.[11] The court held that only a sustained or systematic failure of the board to exercise oversight—such as an utter failure to attempt to assure a reasonable information and reporting system exists—will establish the lack of good faith that is a necessary condition to liability.[12]

The Delaware Supreme Court sharpened *Caremark* in *Marchand v. Barnhill*, holding that when a compliance issue is essential and mission critical, the board's oversight obligation must be more rigorously exercised.[13] The Court found that Blue Bell's board had no committee addressing food safety, no regular process requiring management to keep the board apprised of food safety compliance, and no schedule for the board to consider food safety risks on a regular basis.[14] On those facts, the Delaware Supreme Court reversed the Court of Chancery's dismissal of the derivative complaint, holding that plaintiffs had adequately stated a *Caremark* claim under prong one—the utter failure to implement any reporting system for a mission-critical compliance risk.[15] For companies where AI is core to the product, the revenue engine, or safety-critical decisions, the *Marchand* standard may apply with similar force.

In re Boeing Company Derivative Litigation applied this framework in the aviation safety context. The court found that no board committee had been specifically tasked with overseeing airplane safety, and that the committee charters were silent on the topic.[16] The court further noted that management's safety reports to the board were discretionary rather than systematic.[17] The court concluded that the board had failed to implement any reporting or information systems or controls for safety oversight.[18] The court accordingly denied

---

[10] See *Firemen's Ret. Sys. of St. Louis v. Sorenson*, No. 2019-0965-LWW, 2021 WL 4593777 (Del. Ch. Oct. 5, 2021).
[11] *In re Caremark Int'l Inc. Derivative Litig.*, 698 A.2d 959, 970 (Del. Ch. 1996).
[12] *Id.*
[13] *Marchand v. Barnhill*, 212 A.3d 805, 822 (Del. 2019).
[14] *Id.* at 822–23.
[15] *Id.* at 809, 822–24.
[16] *In re Boeing*, 2021 WL 4059934, at *5.
[17] *Id.* at *29.
[18] *Id.* at *17.

Boeing's motion to dismiss, holding that the plaintiffs' allegations were sufficient to support a reasonable inference that the directors had acted in bad faith by failing to establish any monitoring system for the company's central safety risk.[19] The factors the court considered—whether a committee existed, whether reporting protocols were in place, whether escalation mechanisms were established, and whether the topic appeared on regular board agendas—provide a useful framework for boards evaluating the adequacy of their own AI governance structures.

The 2023 *McDonald's* decision extended this framework to officers, holding that corporate officers owe a fiduciary duty of oversight as to matters within their areas of responsibility.[20] Under *McDonald's*, officers are optimally positioned to identify red flags and either address them or report upward to more senior officers or to the board.[21] Chief Technology Officers, Chief AI Officers, and Chief Compliance Officers who oversee AI deployment should be aware that delegation alone may not be sufficient; under *McDonald's*, they may owe their own independent duty to build information systems and report red flags—bias incidents, model failures, regulatory inquiries—to the board.[22]

Criminal and Regulatory Liability: A Forward-Looking Consideration

The Responsible Corporate Officer (RCO) doctrine, rooted in the Supreme Court's 1943 decision in *United States v. Dotterweich* and developed further in *United States v. Park*, creates strict or near-strict criminal liability for officers in a "responsible relation" to public welfare offenses, even without personal knowledge of the specific conduct at issue.[23] The doctrine has historically been confined to a narrow set of public welfare statutes, principally the Federal Food, Drug, and Cosmetic Act and environmental law. No court has applied it in the AI context, and its extension to AI governance would require either new legislation or a significant doctrinal expansion.

That said, the trajectory of AI regulation suggests this is worth watching. If Congress or federal agencies create AI-specific regulatory obligations with strict liability provisions—a plausible development given the EU AI Act's approach, Colorado's SB 24-205 framework, and proposals at the federal level—the RCO doctrine could become directly relevant to officers overseeing AI deployment. Chancellor Allen himself noted in *Caremark* the increasing tendency, especially under federal law, to employ criminal law to assure corporate compliance with external legal requirements, including environmental, financial, employee and product safety regulations.[24] Should AI-specific regulatory obligations crystallize, officers in a responsible relation to AI systems that touch health care, consumer safety, or financial

[19] *Id.* at *30.
[20] *In re McDonald's Corp. S'holder Derivative Litig.*, No. 2021-0324-JTL, 2023 WL 387292 (Del. Ch. Jan. 26, 2023).
[21] *Id.* at *15.
[22] *Id.* at *15–17.
[23] *United States v. Dotterweich*, 320 U.S. 277, 281 (1943); *United States v. Park*, 421 U.S. 658 (1975).
[24] *Caremark*, 698 A.2d at 969.

services may find themselves evaluated under a similar standard. The prudent course is to build governance systems now that would satisfy such scrutiny if it materializes.

Practical Recommendations

No *Caremark* claim has yet succeeded on AI governance grounds specifically. The recommendations below are best understood not as responses to a demonstrated wave of litigation, but as prudent risk management—and, critically, as the infrastructure that would allow a board to dispose of a *Caremark* claim at the earliest stages should one be brought. This is arguably the most practical reason to act: the value of governance infrastructure is not merely that it reduces the probability of a substantive loss, but that it provides the factual record a court needs to dismiss a derivative claim at the pleading stage, before the company incurs the significant expense of discovery and trial.

The *Marriott* litigation illustrates this dynamic. There, the court dismissed a *Caremark* claim because Marriott's board had implemented reporting systems: it ranked cybersecurity as a top risk, received regular briefings, and engaged outside consultants to audit its practices.[25] Marriott's cybersecurity governance program was, of course, built for a different category of risk. AI risk is less mature from a regulatory standpoint, the standards are less settled, and the threat vectors are more diffuse. The *Marriott* template does not translate to AI governance without adaptation. But the structural features the court found dispositive—committee oversight, regular reporting, third-party audits, and a documented record of board engagement—are transferable as a governance framework regardless of the underlying risk domain.

Translating the *Marchand* and *Boeing* framework into actionable AI governance, boards may wish to consider the following measures: designating a committee with explicit AI oversight responsibility; establishing regular reporting protocols for AI risks, incidents, and compliance, rather than relying on ad hoc management reports; defining what constitutes an AI "red flag"—bias incidents, model failures, regulatory inquiries, vendor audit failures—and requiring escalation; engaging independent third-party auditors rather than relying solely on management's assurances; and ensuring officer-level accountability, so that the CTO, CAIO, and CCO understand their potential fiduciary obligations within their respective domains.[26]

Conclusion

AI governance touches product safety, data privacy, consumer protection, and criminal law simultaneously. The cases discussed above illustrate a consistent judicial expectation: where a risk is central to a company's operations or legal obligations, boards should develop governance structures reasonably designed to monitor it. Companies that are deploying AI at

[25] *Firemen's Ret. Sys.*, 2021 WL 4593777, at *13.

[26] See *Marchand*, 212 A.3d at 822–23; *In re Boeing*, 2021 WL 4059934, at *17–20.

scale would be well served to evaluate whether their current oversight frameworks are commensurate with the risks involved. As the Delaware Supreme Court put it in *Marchand*: "If *Caremark* means anything, it is that a corporate board must make a good faith effort to exercise its duty of care. A failure to make that effort constitutes a breach of the duty of loyalty."[27]

[27] *Marchand*, 212 A.3d at 824.