LAW360

Portfolio Media. Inc. | 230 Park Avenue, 7th Floor | New York, NY 10169 | www.law360.com
Phone: +1 646 783 7100 | Fax: +1 646 783 7161 | customerservice@law360.com

Operational AI Washing: The Next Frontier Of Fiduciary Risk

By **Eric Coleman, Donnie King and Reginald Janvier** (June 9, 2026, 4:16 PM EDT)


Eric Coleman


Donnie King

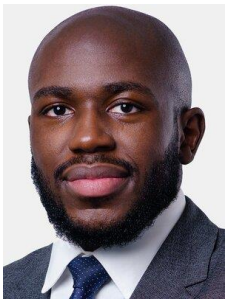
Reginald Janvier

This series has built a unified framework for understanding and defending against operational AI washing claims. The prior four installments traced the liability arc from AI washing disclosure claims through Section 220 demands, the Private Securities Litigation Reform Act's mixed statement trap and motion to dismiss defense strategy.

Each installment rested on the foundation that the board record is simultaneously the plaintiff's primary instrument of attack and the company's primary shield.

The thread connecting each of the first four installments is that the same board records plaintiffs seek under Section 220 of the Delaware General Corporation Law are the records that either shield or expose a company in Rule 10b-5 litigation. That insight now points in a different direction.

This installment asks the upstream question: What does the board owe in the first instance, before disclosures are made or a Section 220 books and records demand letter ever arrives? The answer lies in Caremark's familiar demand for good faith monitoring of risks central to the enterprise.

The Evolution of Caremark in the Artificial Intelligence Era

Under Delaware law, directors have a duty to act in good faith to ensure that reasonable corporate reporting systems exist and that they actively monitor them.

While claims brought under the Delaware Court of Chancery's seminal 1996 ruling in In re: Caremark International Inc. Derivative Litigation are notoriously difficult to litigate, decisions like those in **2019's** Marchand v. Barnhill (food safety) and 2021's In re: Boeing Co. Derivative Litigation (aviation safety) have clarified that when a risk is mission-critical, boards cannot remain passive.

Those cases underscore a consistent theme. When a risk goes to the heart of how the company operates or generates value, the absence of structured board-level oversight itself becomes the alleged breach.

It is important to acknowledge directly that there are still no final, resolved Delaware decisions applying Caremark specifically to AI governance failures. That gap is analytically significant, but it does not leave

boards or practitioners without guidance.

The framework established in Marchand and Boeing is sufficiently developed such that when AI-specific Caremark claims do arrive — the conditions for them are already present in companies deploying AI in regulated, revenue-generating or safety-critical contexts — Delaware courts will apply these existing principles without modification.

The question for boards and their counsel is not whether the doctrine applies, but whether their governance architecture will satisfy it when the first cases are decided. Boards that wait for AI-specific precedent to build their oversight structures will have waited too long.

Assessing the Algorithmic Risk Landscape

Effective AI oversight begins not merely with identifying risk, but with assessing which of those risks are most likely to trigger fiduciary scrutiny.

As a threshold matter, boards should require management to map out how the company is using AI across the business — specifically, whether AI supports internal efficiency, informs regulated decisions, interacts directly with customers or operates in safety-critical environments.

This mapping exercise does more than inform risk management. As Parts 2 and 3 of this series demonstrate, the board record — including how management characterized AI deployment status to the board — is the first document a plaintiff examines in a Section 220 demand and the primary material a federal court evaluates on a motion to dismiss a Rule 10b-5 claim.

The enterprise-level inventory established by a disciplined mapping exercise is therefore the factual baseline against which public disclosures will be measured. A board that has required management to map AI deployment in real time will have records reflecting a good faith, contemporaneous assessment of operational status, rather than a post hoc reconstruction assembled after the complaint is filed.

Once that baseline is established, boards should assess AI risk across several recurring dimensions that implicate regulatory exposure, disclosure obligations and consumer harm:

- Algorithmic bias: Discriminatory outcomes in lending, hiring or underwriting that invite regulatory enforcement and civil liability;

- Data integrity and privacy: Improper data sourcing or violations of privacy laws — the General Data Protection Regulation or the California Consumer Privacy Act — that expose the firm to substantial fines and enforcement actions;

- Model reliability: Hallucinations or unreliable outputs in high-stakes environments such as medicine, finance or compliance decision-making; and

- AI washing: Overstating AI capabilities in public disclosures in a manner that creates direct exposure to securities litigation and derivative claims.

The inclusion of AI washing in this risk classification is deliberate. It situates disclosure risk as a governance issue, not merely a communications failure. When board materials reflect a deployment gap that public disclosures do not acknowledge, the board itself becomes a potential defendant in derivative litigation, not merely a witness to it.

Boards should also consider the question of director fluency. Caremark does not require directors to become engineers or data scientists. But, at the same time, the business judgment rule does not protect directors who rubber-stamp management's technical assessments without probing whether those assessments are complete, consistent with the public record and free of known red flags.

However, it does require sufficient fluency to evaluate management's risk assessments and recognize red flags when they arise. In practice, that means directors should be able to do three things:

- Evaluate whether management's characterization of AI deployment status is consistent with the technical record;
- Identify when a chief technology officer's cautionary assessment diverges from the narrative in a press release or earnings call script; and
- Ask sufficiently probing questions that the answers are memorialized in board minutes.

That last point is not procedural housekeeping. As Part 2 explains, board minutes are among the first documents produced under a Section 220 demand. The difference between a protective minute entry and an inadequate one is concrete.

A minute that states "The Chief Technology Officer provided an update on the company's AI initiatives; no action was taken" provides no Caremark protection and no scienter defense. A minute that states "The Technology Committee received the chief technology officer's quarterly AI risk assessment. Directors raised questions regarding the model validation backlog in the underwriting system and requested written responses within 30 days; follow-up scheduled for the November meeting" reflects the kind of documented, deliberate engagement that satisfies both prongs of the oversight standard.

Director fluency is the mechanism that produces the second kind of minute, not the first. That distinction matters because, under Caremark, the business judgment rule protects directors' good faith oversight process, not the technical success or failure of the systems they oversee.

What Caremark Requires: Governance Architecture and Liability Boundaries

Because Caremark polices process rather than outcomes, effective AI oversight turns on whether a board has established clear ownership, reliable reporting systems and a defensible response to red flags.

Organizational Ownership and Reporting Systems

Satisfying the good faith requirement of Caremark requires a dual focus on organizational ownership and repeatable process. Boards should first ensure that management has clearly designated responsibility for AI governance, whether through a chief AI officer, a cross-functional AI risk committee or an AI ethics board.

From a litigation perspective, the existence of a named owner matters less than whether that owner produces regular, documented reporting that reaches the board or a delegated committee. From a fiduciary perspective, the board's role is not to manage these functions but to ensure they are empowered, with an adequate budget, technical staff and, critically, direct access to the audit, risk or technology committee.

Boards should also look to their existing cybersecurity governance frameworks as a practical blueprint. Much like the U.S. Securities and Exchange Commission's 2023 cybersecurity disclosure rules, which require public companies to report material cyber incidents, AI oversight demands a systems-based approach to risk identification, escalation and disclosure.

This analogy is especially useful because courts and regulators already understand cyber governance as a Caremark-relevant risk domain. Because AI and cyber risks are frequently intertwined, such as models trained on compromised or improperly sourced data, boards can often integrate AI oversight into established cyber reporting lines. This approach minimizes governance fatigue while leveraging familiar escalation protocols, regularized reporting and whistleblower mechanisms.

There is a direct connection between the reporting structures described here and the Section 220(b)(3) incorporation-by-reference doctrine addressed in Part 2. When a board conditions production of its records on the plaintiff's acceptance that those records will be deemed incorporated by reference into any subsequent complaint, the doctrine is only a defense if the records are clean.

A governance architecture that produces regular, documented chief technology officer assessments — reflecting actual deployment status, known constraints and good faith risk estimates — creates exactly the kind of internal record that, once incorporated, defeats rather than confirms a plaintiff's theory.

The reporting structure, in other words, does not merely satisfy Caremark. It operationalizes the most powerful prelitigation defense available under current Delaware law.

Process as the Liability Boundary

It is equally important to underscore what Caremark does not require. Delaware law does not impose liability for flawed technology, failed investments or imperfect outcomes. Its focus remains squarely on process.

Caremark liability remains an exacting standard, triggered only by a sustained or systemic failure of oversight amounting to bad faith, such as the absence of any reporting system or a conscious disregard of known red flags. This distinction will be critical in AI-related litigation, where plaintiffs may attempt to repackage technical underperformance as fiduciary misconduct.

Practical Oversight Habits

To remain within this protective boundary, boards overseeing AI should adopt several concrete practices.

Clarify responsibility.

Explicitly assign AI oversight to a specific board committee, and consider which committee is the right fit. For companies where AI drives revenue or customer-facing decisions, the risk or technology committee is the natural home, because those committees are already structured to receive operational and technical escalations. For companies where AI's primary footprint is in financial reporting or compliance, the audit committee may be better positioned.

The answer matters less than the clarity. A board that has not designated a committee owner has, from

a Caremark standpoint, no one watching. The designation should appear in the committee's charter, not merely in board minutes, so that it survives committee reconstitution and is visible to any plaintiff conducting a Section 220 review of the company's governance documents.

Document engagement.

Ensure board minutes reflect substantive discussion and meaningful oversight questions regarding model validation, limitations and stress testing.

As Part 2 explains, minutes that reflect only conclusions — without questions asked, answers received or follow-up directed — provide no protection under Section 220(b)(3) and no defense to a scienter argument at the motion to dismiss stage.

Audit the algorithm.

Integrate AI risks into existing internal audit, compliance and risk management functions.

At a minimum, the AI audit scope should cover model validation protocols, data sourcing documentation, output monitoring for accuracy and bias, and the gap between AI capability representations in public disclosures and actual deployment status.

The board should establish, in writing, what findings trigger escalation from the committee to the full board, and that threshold should be lower than the threshold for other operational risks, given the relative novelty of AI governance and the pace at which deployment gaps can develop.

Review disclosures.

Confirm that marketing materials and SEC filings describing AI capabilities are grounded in technical reality.

As Part 3 establishes, a single mixed statement in an earnings call or press release — one that blends a present-fact claim about AI maturity with a forward-looking projection — can defeat the PSLRA safe harbor and render months of careful governance work irrelevant. The board's review obligation runs to every public channel, not just the Form 8-K. The governance mechanism for that review should be specified and documented.

In practice, this means assigning a designated owner to cross-reference every external AI capability claim against the most recent internal technical assessment before the statement is made public.

Conclusion

Across five installments, this series has traced a single arc: that operational AI washing exposure arises under specific conditions. These include AI being embedded in core operations, a narrative drift between internal reality and public explanation, and a board record that either supports or contradicts the company's public disclosures.

These same conditions ultimately determine whether a board has fulfilled its Caremark duty of oversight.

The half-truth doctrine described in Part 1, the Section 220 demand mechanics of Part 2, the mixed statement trap of Part 3 and the motion to dismiss architecture of Part 4 all resolve to the same upstream question: Did the board build and maintain a governance structure that reflects operational reality, and does the public record say the same thing the board record says?

When the answer is yes, the company is positioned to win at every stage. When the answer is no, no amount of litigation strategy, however skillfully executed, fully closes the gap.

Eric D. Coleman and Donnie M. King are partners, and Reginald Janvier is an associate, at Akerman LLP.