HANDOUT 2

AI Risk Disclosure and the Regulator in the Room: What Public Company Directors and Senior Executive Teams Need to Know About the SEC

Board AI Governance Training Session | IterateOn | April 2026

Purpose of this handout

This document serves as your legal grounding for the April 13 session. It explains what regulators — especially the SEC — are already requiring from public company boards on AI. It explains what courts have already decided. And it explains what can go wrong, personally and for your company, if AI governance is treated as a future problem rather than a current one.

The risks described here are real; they are already showing up in enforcement actions and court rulings, and they apply to you whether you sit on the board or report to it.

This document is not legal advice. The authors are not your attorneys. Nothing here should be read as a substitute for advice from qualified legal counsel familiar with your specific situation.

PART 1: THE REGULATORY BASELINE

1.1 The SEC Has Not Waited for Congress

AI-specific legislation at the federal level remains fragmented and incomplete. But the absence of a dedicated AI statute does not mean the SEC is watching from the sidelines. The Commission has moved aggressively under existing authority, using the same disclosure framework that governs cybersecurity, material risk, and management's discussion and analysis — applying it now to artificial intelligence.

Directors should understand the existing architecture before the session, because it explains why AI governance is already a board-level legal matter, not a future concern.

The Materiality Standard Applies Now

Under longstanding SEC doctrine, companies must disclose information that a reasonable investor would consider important in making an investment decision. The SEC has made clear — through guidance, comment letters, and enforcement — that AI meets this threshold in many contexts.

In June 2024, **Eric Gerding, Director of the SEC's Division of Corporation Finance**, publicly identified AI as a disclosure priority, noting that the Division was observing a significant increase in companies mentioning AI in annual reports and would scrutinize whether those disclosures were tailored and substantive, or vague and boilerplate.

The Cybersecurity Disclosure Rule: The Direct Precedent

In September 2023, the SEC's cybersecurity disclosure rules became effective (Regulation S-K Item 106; Form 8-K Item 1.05). These rules require public companies to:

- Disclose material cybersecurity incidents within four business days of determining materiality (Form 8-K, Item 1.05)
- Describe annually their processes for identifying, assessing, and managing material cybersecurity risks (Form 10-K, Regulation S-K Item 106)
- Describe the board's oversight of cybersecurity risk, including which committee is responsible and how the board is kept informed
- Describe management's role in assessing and managing cybersecurity threats

Why does this matter for AI? Because the SEC's own Investor Advisory Committee and senior staff have explicitly framed AI disclosure as the next chapter in this story. The cybersecurity rules are the template. The AI rules — whether formal or guidance-based — are being written now.

PART 2: WHAT THE SEC IS SAYING ABOUT AI, SPECIFICALLY

2.1 The IAC Recommendations (December 2025)

At its December 4, 2025 meeting, the SEC's Investor Advisory Committee voted to recommend that the Commission issue formal guidance on AI-related disclosure. The recommendations are specific and instructive for any board trying to understand what "adequate" AI governance disclosure looks like:

- Issuers should **define what they mean by artificial intelligence** — generic references to "AI" in risk factors are insufficient
- Issuers should **disclose whether the board or a board committee oversees AI deployment** — and if not, why not
- Disclosures should cover AI's impact on **internal operations**: workforce changes, financial reporting implications, governance structures, and cybersecurity risk
- Disclosures should cover AI's role in **products and consumer-facing services**: investment levels, integration depth, and regulatory exposure
- Guidance should integrate into existing Regulation S-K items (101, 103, 106, and 303) rather than create a new subchapter — meaning the framework is already largely in place

The IAC cited research findings that only 40% of S&P 500 companies provide any AI-related disclosures, and only 15% disclose information about board oversight of AI — despite 60% of S&P 500 companies identifying AI as a material risk. That gap is where the liability lives.

2.2 The Current Commission's Posture

Chair Paul Atkins, who took office in 2025, has publicly stated that the SEC's existing principles-based rules are sufficient to address AI disclosure — and that prescriptive checklists are not the answer. His position: materiality is the test, and companies that are using AI in material ways are already required to say so under current rules.

This posture has a specific implication for directors: the absence of a formal AI disclosure rule does not provide cover. If AI is material to your business — your operations, your risk profile, your competitive positioning — existing SEC rules already require you to disclose it. The question is not whether the rules apply. It is whether your disclosures are adequate.

2.3 Enforcement Is Already Happening

The SEC has not waited for new rules to act. In 2024, it took several enforcement actions against companies for misrepresenting their AI capabilities to investors — a practice the Commission calls "AI washing." Key examples:

- Two investment advisory firms were charged for falsely claiming AI drove their investment decision-making when it did not
- In January 2025, the SEC reached a non-monetary settlement with a consumer technology company for making false and misleading statements about core aspects of its AI product
- The SEC's Division of Examinations listed AI as an examination priority in both its 2025 and 2026 priority announcements
- In February 2025, the SEC launched the Cyber and Emerging Technologies Unit (CETU) — a dedicated enforcement unit focused on AI and technology-related misconduct

The enforcement pattern is instructive: the SEC is not only concerned with companies that overstate their AI capabilities. It is also concerned with companies that fail to disclose AI-related risks, limitations, and governance structures that would be material to investors.

PART 3: THE BOARD'S EXPOSURE

3.1 The Caremark Doctrine and AI Oversight

Under the Caremark standard (*In re Caremark Int'l Inc. Derivative Litig., 698 A.2d 959 (Del. Ch. 1996*)), boards face personal liability when they fail to implement and monitor systems for overseeing mission-critical risks. Delaware courts have refined this doctrine through subsequent cases — most relevantly *Marchand v. Barnhill (212 A.3d 805 (Del. 2019))* — to hold that for risks that go to the core of a company's business, board oversight must be more rigorously exercised.

Caremark applies to both public and private companies incorporated in Delaware. For nonprofit boards, the doctrine does not apply in the same technical form — there are no shareholders to bring a derivative suit — but nonprofit directors are not insulated from oversight liability. State attorneys general can bring claims against nonprofit boards without meeting the stringent bad faith standard required in corporate derivative suits. The IRS can impose excise taxes on nonprofit board members for certain breaches of duty. And as AI becomes central to how nonprofits deliver services and manage data, the adequacy of board oversight will be scrutinized under whatever state law framework applies.

AI is rapidly moving into mission-critical territory for many companies. It is influencing pricing, lending decisions, hiring, customer service, compliance functions, and financial reporting. As

that happens, the argument that a board had no responsibility to establish oversight structures becomes harder to sustain.

There are two Caremark failure modes boards should understand:

- **Failure to implement oversight systems.** The board never established a mechanism for receiving information about AI risk, oversight structures, or deployment decisions.
- **Failure to act on red flags.** The board received warnings — in management reports, audit findings, or press coverage — and failed to respond.

While Delaware courts have generally set a high bar for Caremark claims, the cybersecurity analogy is instructive: derivative suits were filed against boards of Marriott, SolarWinds, and others after major incidents. Most were dismissed — but not without significant legal cost, management distraction, and reputational consequence. And courts have made clear that the analysis evolves as risks become more widely understood by directors.

3.2 Disclosure Accuracy Is a Director's Responsibility

There is a second, more direct liability path: the accuracy of public disclosures. If a company discloses that it has "robust AI governance structures" or that "the board actively oversees AI risk" — but that language does not reflect actual practice — directors face exposure under Section 10(b) of the Securities Exchange Act and SEC Rule 10b-5 for materially misleading statements.

This is not hypothetical. Legal scholars have argued — and some courts have begun to consider — that Caremark liability can attach specifically to director oversight failures around disclosure accuracy, not just operational failures. A board that approves AI-related disclosures it has not meaningfully reviewed is exposed.

The key question for every director to answer honestly:

If your company's 10-K states that the board or a committee oversees AI risk, what does that oversight actually consist of? How often does it occur? What information does the board receive, from whom, and in what form? Could you describe it under oath?

If the answer is unclear, the disclosure may be a liability, not a protection.

3.3 The Privilege Problem: A Court Has Already Ruled

There is a third exposure area that most boards have not yet registered, and it landed in federal court just weeks ago.

On February 10, 2026, Judge Jed S. Rakoff of the U.S. District Court for the Southern District of New York ruled in *United States v. Heppner* (No. 25-cr-00503-JSR) that documents a defendant created using a public AI platform were not protected by attorney-client privilege and could be handed directly to federal prosecutors. His written opinion followed on February 17. The court itself described it as a question of first impression nationwide — meaning no federal court had ruled on it before.

The facts matter because they are not exotic. Bradley Heppner, a financial services executive facing securities fraud and wire fraud charges, used Claude — the publicly available consumer version — to research legal questions related to his case. He typed in information he had received from his own attorneys. He later shared the AI outputs with his legal team. He believed this was part of his legal defense preparation. The government disagreed. So did Judge Rakoff.

Public AI vs. Private AI: the distinction that governs your legal exposure

Private AI means the company controls all three things that matter: the hardware it runs on, the model itself, and the data that flows through it. What goes in stays in.

Public AI is the opposite. With public AI — consumer-facing tools like ChatGPT, Claude.ai, Google Gemini, or Microsoft Copilot — you control none of those three things. Your input goes to the AI company's servers, runs on their infrastructure, and is governed by their terms of service, which typically allow data collection, storage, and in some cases disclosure to "governmental regulatory authorities."

The *Heppner* ruling applies specifically to public AI. But most employees at most companies are using public AI every day — often without realizing what that means for confidentiality or legal privilege.

The court's reasoning was blunt and has three parts that every board member and executive should understand.

First: an AI tool is not a lawyer. The attorney-client privilege applies to communications between a client and an attorney. Claude is not an attorney. It has no law license, owes no duty of loyalty, and cannot form an attorney-client relationship. Judge Rakoff wrote that this fact alone "disposes of Heppner's claim of privilege."

Second: public AI is not confidential. Privilege requires that a communication was intended to be, and actually was, kept confidential. When you type something into a public AI platform, you are sharing it with a third-party company whose privacy policy does not guarantee confidentiality. Anthropic's own terms, cited by the government in the case, permit data collection and disclosure to governmental authorities. The moment Heppner typed his attorneys' advice into Claude, he handed it to a third party — and the privilege was gone.

Third: you cannot fix it after the fact. Heppner's lawyers argued that because he eventually shared the AI outputs with them, the materials became privileged. The court rejected this. Privilege must exist at the moment of the communication. Routing something through your attorney after the fact does not restore a privilege that was never there.

Fourth: one prompt may have stripped privilege from the underlying lawyer communications, too. Heppner did not just ask Claude generic questions. He typed in the information his attorneys had given him — their legal strategy, their assessment of the case — directly into a public AI platform. The court agreed with the government that, by doing so, he waived the privilege over those original attorney-client communications. By sharing what his lawyer told him with a third party, he potentially gave the government access not only to the AI documents but also to the privileged legal advice that went into them.

Courts have not yet resolved how broadly that waiver extends. Under the **subject matter waiver doctrine**, disclosing some privileged communications on a topic can strip privilege from all related communications on the same subject. Whether typing your lawyer's advice into a public AI tool triggers that broader waiver — effectively **unprotecting an entire category of attorney-client conversations** — is a question that is now moving through the courts. The answer may be yes.

The court also rejected work **product protection**. The work product doctrine shields materials prepared by or at the direction of counsel in anticipation of litigation. Because Heppner created the AI documents on his own initiative — not at his lawyers' direction — that protection did not apply either. The court noted, however, that the analysis might have been different if his counsel had directed him to use the tool. That distinction matters for how companies structure AI use going forward.

The immediate implications for board members and senior executives are concrete:

- Any employee who uses a public AI tool to analyze a legal question, evaluate liability, draft a response to a regulator, or research a compliance issue may be creating a

discoverable record that opponents can obtain.

- This applies in civil litigation, workplace investigations, regulatory inquiries, and internal business analysis — not just criminal cases.

- It applies to executives as much as to employees. A CEO who types strategic legal questions into ChatGPT before a board meeting has not had a privileged conversation. They have created a document.

- If the employee types what their lawyer told them into a public AI tool, they may have waived privilege over those original attorney-client communications — not just the AI outputs. One careless prompt can potentially strip protection from an entire category of legal advice.

- Sending the AI output to your general counsel afterward does not fix it.

The practical question boards and executive teams need to answer is not whether to use AI for legal and compliance work — it is which type of AI, under what terms, with what oversight. The privilege analysis differs materially across three tiers.

The first tier is public AI: ChatGPT, Claude.ai, Google Gemini, and the free or standard consumer versions of Microsoft Copilot. *Heppner* applies directly here. There is no reasonable expectation of confidentiality. The platform's terms permit data collection, training use, and disclosure to regulatory authorities. Any sensitive legal information typed into these tools is effectively handed to a third party. Privilege is destroyed at the moment of input.

The second tier is enterprise AI with verified contractual confidentiality. The most common example is Microsoft Copilot — but the license alone is not enough. Most companies that "have Copilot" are running it under terms that do not provide the protections they assume. What is actually required: a Microsoft 365 enterprise license at the appropriate tier, a signed Data Protection Addendum, and a confirmed zero-training configuration — meaning your IT team has verified, not assumed, that Microsoft is not using your prompts and responses to train or improve its models. All three conditions must be met and auditable. If your general counsel cannot point to documentation confirming each one, you do not have the protection this tier describes.

The same logic applies to other enterprise deployments: private cloud arrangements, vendor agreements with explicit prohibitions on training data use, and similar setups where the company has committed contractual terms in writing — not just relied on default settings.

This is a common situation: the majority of companies with Copilot or other enterprise AI tools have not done all three of these things. They signed up for the license, perhaps accepted a vendor agreement, and assumed the rest was handled. In most cases, it was not. The signed

DPA and verified zero-training configuration require deliberate action by legal, IT, and procurement working together — and that coordination rarely happens spontaneously. This is exactly the kind of gap that surfaces when litigation begins, and someone starts asking what the actual configuration was.

There is also a gap that legal agreements alone cannot close. Even when the contract is right — the DPA is signed, and the zero-training configuration is confirmed — data can still leak if the underlying infrastructure is shared. Microsoft's enterprise terms provide legal assurances, but those assurances describe what Microsoft will do with your data, not what the architecture prevents. On shared compute infrastructure, data from one tenant can, in some configurations, be exposed to another through misconfigurations, vulnerabilities, or side-channel risks that no contract anticipates. Microsoft's indemnification provisions do not fully compensate for the business, competitive, or legal harm that follows a data exposure — they address Microsoft's liability, not yours. A signed DPA tells you what recourse you have after something goes wrong. It does not guarantee that nothing will.

Setting aside the technical risks for a moment, even in the best-case scenario — where the contract is right, the configuration is verified, and no data has leaked — one limitation survives that no legal agreement can fix: the AI is still not an attorney."There is no third-party disclosure problem. That is meaningfully better than public AI. But one limitation survives regardless of how good the contract is: the AI is still not an attorney. Conversations with an enterprise AI tool, standing alone, are not privileged communications. The work product and Kovel doctrines may offer protection when the tool is deployed at counsel's direction as part of legal strategy work — but that protection is not automatic, and courts have not yet ruled definitively on its boundaries.

The third tier is fully private AI: a model the company owns, running on hardware the company controls, with no external data flow whatsoever. This is the strongest position available today. The confidentiality problem is solved — there is no third party, no vendor terms to worry about, and no disclosure risk. But the attorney limitation remains regardless. An AI system is not a lawyer, no matter who owns the hardware. Conversations with a private AI are not privileged, standing alone. **Where they can be protected is under the work product doctrine**, if the tool is deployed at counsel's direction, is used to assist with specific legal strategy work, and that relationship is documented. That requires intentional structuring — not just private infrastructure.

The distinction between these three tiers needs to be policy, not assumption — and it needs to reach every employee who uses AI for anything that touches legal, compliance, or regulatory matters, which at most companies today means nearly everyone.

PART 4: PRACTICAL IMPLICATIONS FOR DIRECTORS

4.1 AI Is Not Yet a Separate Disclosure Category — But Treat It Like One

The SEC has not yet adopted a standalone AI disclosure rule. But the IAC recommendations, staff guidance, and enforcement record together create a de facto standard that any reasonable board should follow. The integration of AI disclosure into existing Regulation S-K items (particularly Items 101, 103, 106, and 303) means the infrastructure is already there. The gap is in execution.

Directors should expect the following questions from regulators, plaintiffs' attorneys, and sophisticated investors — and should be able to answer them based on actual board activity, not drafted disclosure language:

- How has the board defined what AI means for purposes of oversight and disclosure at this company?
- Which committee, if any, is responsible for AI oversight? What is its charter and meeting frequency?
- What information does the board receive about AI deployment decisions, model behavior, and risk exposure?
- How does the company assess whether an AI-related development is material and requires disclosure?
- Are the company's AI-related disclosures reviewed by legal counsel and the board before filing?
- Has the board assessed third-party AI dependencies — including model providers whose behavior the company does not control?

4.2 The AI Exposure Surface Is Larger Than Cybersecurity

In the cybersecurity era, the board's oversight obligation was primarily reactive: was the company managing breach risk and incident response? AI introduces a fundamentally different oversight challenge because the risk is generative and operational, not just defensive.

AI systems make and influence decisions — in pricing, credit, hiring, compliance, customer communication, and increasingly in financial reporting. This means the exposure surface is not limited to a security incident. It extends to:

- **Algorithmic bias claims** — particularly in employment, lending, and consumer-facing applications subject to anti-discrimination law

- **Model drift and error accumulation** — where an AI system's outputs degrade over time in ways that are not immediately visible to management. This risk is accelerating: Google's Nested Learning paradigm (published at NeurIPS 2025) is designed to make continuous in-model learning the default, meaning AI systems will increasingly update their own behavior as they operate. Governance frameworks built around periodic audits will not keep pace.

- **Vendor dependency and data leakage** — where institutional knowledge is embedded in a third-party model the company does not control

- **Autonomous decision liability** — where AI acts without a human decision-maker in the loop, and accountability becomes legally ambiguous

Each of these creates disclosure obligations, potential litigation exposure, and — where they touch mission-critical operations — Caremark-style oversight requirements.

PART 5: SIX QUESTIONS EVERY BOARD SHOULD BE ABLE TO ANSWER

These are the six questions your board should be able to answer. You may hear more good questions during the session too. If your company cannot answer them clearly and specifically — not in boilerplate — that gap is itself a governance finding.

1	What AI systems does this company currently operate that could be considered material to investors, operations, or risk exposure?
2	Has the board formally designated a committee responsible for AI oversight? Is that designation documented and reflected in the committee charter?
3	What information does the board receive about AI risk — how often, in what format, and from whom? Is that process documented?
4	Do our current 10-K disclosures accurately describe the board's actual AI oversight practices? Would those descriptions hold up to a plaintiff's attorney's scrutiny?
5	Has legal counsel reviewed our AI-related risk factor disclosures in light of the IAC recommendations and SEC staff guidance from 2024–2025?

6	Have we assessed our third-party AI dependencies — including which vendor models we rely on, what data we are providing to them, and whether that creates training data or competitive risk?

A Note on the Pace of Change

The regulatory picture described in this handout reflects the environment as of early 2026. It will change. The SEC may issue formal AI disclosure guidance. The IAC recommendations may be adopted in whole or in part. Enforcement actions will accumulate. State-level AI regulation — particularly in California — is already adding additional obligations for some companies.

The appropriate board posture is not to wait for formal rules to arrive. It is to build governance structures and disclosure practices now that will hold up to scrutiny under whatever framework emerges. The companies that do this early will be better positioned legally, competitively, and reputationally.

The companies that wait are betting that the SEC, plaintiffs' attorneys, and institutional investors will not ask hard questions about AI governance before the rules are fully formed. Based on the trajectory of the past two years, that is not a safe bet.

A final note on scope: the SEC framework described in this handout is one layer of a rapidly expanding legal landscape. Our session on AI Governance for Board Directors addresses how EU, federal, and state AI laws — including the EU AI Act, the Colorado AI Act, and emerging automated decision-making regulations — interact with and extend beyond the disclosure obligations described here. The governance logic in this handout applies across that broader environment, and directors and executive teams should expect legal exposure to expand well beyond SEC disclosure as AI law matures across jurisdictions.

About This Paper

This paper was prepared by Iterate's executive team, including research drawn from Claude (Anthropic) and cross-referenced by Iterate's Private AI Generate platform. Authors were primarily Magnus Tagtstrom, Corporate VP, who worked on GDPR-compliant products in Europe for Circle K, Dave Jenkins, VP of Research, who ran OEM partnerships for Red Hat in Asia and EMEA, and Jon Nordmark, CEO, who served on Colorado's governor- and legislature-appointed AI Task Force as Colorado worked through the first broad-sweeping AI law passed in the United States.

References

SEC Division of Corporation Finance, Statement on AI Disclosure Priorities, Eric Gerding, June 24, 2024

SEC Cybersecurity Risk Management, Strategy, Governance, and Incident Disclosure Rules, Release No. 33-11216, effective September 5, 2023 (Reg. S-K Item 106; Form 8-K Item 1.05)

SEC Investor Advisory Committee, "Disclosure of Artificial Intelligence's Impact on Operations," December 4, 2025

SEC Division of Examinations, 2025 Examination Priorities (October 2024); 2026 Examination Priorities (November 2025)

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

Marchand v. Barnhill, 212 A.3d 805 (Del. 2019)

SEC Enforcement Actions re: AI Washing (2024–2025), including non-monetary settlement with consumer technology company (January 2025)

SEC Cyber and Emerging Technologies Unit (CETU), established February 2025

United States v. Heppner, No. 25-cr-00503-JSR (S.D.N.Y.), oral ruling February 10, 2026; written opinion February 17, 2026 (Judge Jed S. Rakoff) — AI-generated documents not protected by attorney-client privilege or work product doctrine