A FIELD GUIDE FOR BOARDS, AUDIT COMMITTEES & GENERAL COUNSEL

# AI Risk Disclosure and the Regulator in the Room.

What public company directors and senior executive teams need to know about the SEC — and why "we'll address it later" is no longer a defensible posture.

- SEC REG S-K
- CAREMARK / MARCHAND
- UNITED STATES V. HEPPNER
- IAC RECOMMENDATIONS DEC 2025

AUTHORS

**Magnus Tagtstrom · Dave Jenkins · Jon Nordmark**

Iterate.ai Executive Team

EDITION

**Handout 02**

Board AI Governance training

SUBJECT

**SEC Disclosure & Privilege**

Director & senior-executive exposure

ABOUT THIS PAPER

# This paper is your legal grounding for the conversations boards are about to have — whether they want to or not.

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 the company, when AI governance is treated as a future problem rather than a current one.

The risks described here are real. They are showing up in enforcement actions, examination priorities, and court rulings issued within the last quarter. They apply whether you sit on the board or report to it.

NOT LEGAL ADVICE

The authors are not your attorneys. Nothing in this document should be read as a substitute for advice from qualified legal counsel familiar with your specific situation. Laws and rules in this area are evolving rapidly; the snapshot here reflects publicly available information as of April 2026.

WRITTEN FOR

Public company directors, audit committee chairs, CFOs, general counsel, and senior executives whose disclosures, decisions, or signatures touch AI.

PUBLISHED BY

**Iterate.ai** — San Jose, CA & Denver, CO. Private AI infrastructure for the enterprise.

CONTENTS

# What's inside.

Read it front to back, or jump to the parts your role cares about most. The Heppner case on p.10 is the page most general counsel ask to bookmark.

01 OPENING

# The SEC has not waited for Congress.

AI-specific federal legislation remains fragmented and incomplete. The absence of a dedicated AI statute is being mistaken for an absence of regulation. It is not the same thing.

**The Commission has moved aggressively under existing authority — using the same disclosure framework that already governs cybersecurity, material risk, and MD&A.**

**Under longstanding SEC doctrine, companies must disclose information that a reasonable investor would consider important in making an investment decision. That is the materiality standard, and it has been the spine of US disclosure law for half a century.**

The SEC has now made clear — through guidance, comment letters, and enforcement — that AI clears the materiality bar in many contexts. The rules already apply. The question is whether your disclosures already comply.

In June 2024, Eric Gerding, Director of the SEC's Division of Corporation Finance, publicly identified AI as a disclosure priority. He noted 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.

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

THE ARCHITECTURE

## Existing rules already reach AI.

| | |
|---|---|
| ITEM 101 | Business description — including material AI deployments. |
| ITEM 103 | Legal proceedings — AI-related litigation, enforcement. |
| ITEM 106 | Cyber risk management, strategy, governance. **The template the SEC is extending to AI.** |
| ITEM 303 | MD&A — known trends and uncertainties material to operations. |

*The IAC's December 2025 recommendation explicitly directs that AI disclosure integrate into these existing Reg S-K items — not into a new subchapter.*

STATED PRIORITY · JUNE 2024

**The Division of Corporation Finance is watching whether AI disclosures are *tailored and substantive* — or vague and boilerplate.**

**Eric Gerding**
Director, Corporation Finance

02 THE DIRECT PRECEDENT

# The cybersecurity disclosure rule is the template.

In September 2023, the SEC's cyber disclosure rules became effective. The SEC's own Investor Advisory Committee and senior staff have explicitly framed AI disclosure as the next chapter in this story.

TIMELINE · FROM CYBER TO AI

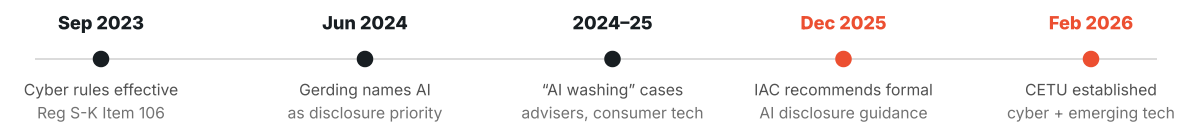


CYBERSECURITY RULE — EFFECTIVE 2023

## What it requires.

- Disclose material cyber incidents within four business days *(Form 8-K Item 1.05)*.
- Describe annually the processes for identifying and managing cyber risks *(Form 10-K Item 106)*.
- Describe the board's oversight — including which committee is responsible and how the board is kept informed.
- Describe management's role in assessing and managing the threats.

AI DISCLOSURE — BEING WRITTEN NOW

## What the parallel looks like.

- Disclose AI-related developments that are material to investors.
- Describe processes for identifying, assessing, and managing AI risks.
- Describe the board's oversight of AI — which committee, what cadence, what information.
- Describe management's role in operating and governing AI systems.

**The rules — whether formal or guidance-based — will rhyme with the cyber rules. Boards that build for one are building for the other.**

03 THE IAC RECOMMENDATIONS · DECEMBER 4, 2025

# What "adequate" AI governance disclosure now looks like.

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. They function as an early answer key.

**01** **Define what you mean by "AI."**

Generic references in risk factors are insufficient. Issuers should specify which systems, capabilities, and use cases are in scope — classical ML, generative models, agents, retrieval-augmented systems — rather than rely on a single undefined term.

**02** **Say whether the board (or a committee) oversees AI — and if not, why not.**

The IAC explicitly recommends "disclose or explain." A blank answer is itself a disclosure.

**03** **Cover AI's impact on internal operations.**

Workforce changes. Financial reporting implications. Governance structures. Cybersecurity risk. These are the categories regulators expect to see addressed substantively in MD&A, not deflected to a single risk-factor paragraph.

**04** **Cover AI's role in products and consumer-facing services.**

Investment levels, integration depth, and regulatory exposure — described in terms specific enough that an investor could meaningfully compare two issuers.

**05** **Integrate into existing Reg S-K items 101, 103, 106, and 303.**

Not a new subchapter. The framework is already largely in place — which means the obligations are already largely in place too.

---

WHERE THE LIABILITY LIVES · IAC RESEARCH

**60%**

of S&P 500 companies identify AI as a **material risk**.

**40%**

provide any AI-related disclosures at all.

**15%**

disclose information about **board oversight** of AI. *That gap is the exposure.*

04 POSTURE & ENFORCEMENT

# The Commission's "principles-based" posture is not cover.

Chair Paul Atkins, who took office in 2025, has stated publicly that 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 material AI use is already disclosable today.

For directors, the implication is specific: **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 — the existing rules already require you to disclose it.

The question is not whether the rules apply. The question is whether your disclosures are adequate. That is a higher bar than it sounds, because "adequate" is now being measured against the IAC recommendations, the cyber-rule analogy, and recent enforcement actions — not against the boilerplate language drafted in 2023.

THE DIRECTOR'S TEST

**If your 10-K says the board oversees AI risk, can you describe what that oversight *actually consists of* — cadence, format, information flow, who reports — under oath?**

ENFORCEMENT IS ALREADY HAPPENING

AI WASHING — 2024

**Two investment advisory firms charged.**

Falsely claimed AI drove their investment decision-making. It did not.

MISLEADING STATEMENTS — Jan 2025

**Non-monetary settlement with consumer tech company.**

False and misleading statements about core aspects of its AI product.

EXAM PRIORITY — 2025 & 2026

**AI listed by Division of Examinations.**

Two priority announcements in a row.

NEW UNIT — Feb 2025

**CETU launched.**

The Cyber and Emerging Technologies Unit — 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. Silence is not a safe harbor.

05 THE CAREMARK DOCTRINE

# When AI becomes mission-critical, oversight becomes personal.

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 the doctrine through subsequent cases — most relevantly *Marchand v. Barnhill (212 A.3d 805, Del. 2019)* — holding that for risks that go to the **core of a company's business**, board oversight must be more rigorously exercised.

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.

## TWO CAREMARK FAILURE MODES

**01**

**Failure to implement.**

The board never established a mechanism for receiving information about AI risk, oversight structures, or deployment decisions.

**02**

**Failure to act on red flags.**

The board received warnings — in management reports, audit findings, or press coverage — and failed to respond.

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

### WHO CAREMARK REACHES

**Both public and private companies incorporated in Delaware.**

If your entity is Delaware-formed, the doctrine reaches you.

### NONPROFITS ARE NOT INSULATED

Nonprofit directors face a different but real oversight landscape:

- State attorneys general can bring claims without the stringent bad-faith standard.
- The IRS can impose excise taxes on individual board members.
- State law oversight standards apply as AI becomes central to service delivery.

**Translation:** "We're a nonprofit, so this doesn't apply" is not a defense. It is a different jurisdictional posture, not a shield.

06 A SECOND, MORE DIRECT LIABILITY PATH

# Disclosure accuracy is a director's responsibility.

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. So is one that relies on outside counsel's drafting language without verifying that the language describes what is actually happening.

The key question for every director to answer honestly is the next paragraph in this paper. It looks innocuous. It is not.

**THE HONEST QUESTION**

**If your company's 10-K states 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.**
>
> *That sentence is the spine of director AI exposure.*

**WHAT AN ADEQUATE ANSWER LOOKS LIKE**

- A named committee with a documented charter that says "AI."
- A stated cadence — quarterly, at minimum, with documented agendas.
- A named executive who reports in, with a written reporting template.
- Minutes and materials that demonstrate the oversight was actually exercised.
- Disclosures drafted to *describe* the above, not aspire to it.

**Note for the audit committee:** the same logic applies to MD&A. "We invest in AI" is not a disclosure — it is an admission of materiality without a description of how it's managed. The SEC has been clear that this is exactly the language it is examining.

07 THE PRIVILEGE PROBLEM · A COURT HAS ALREADY RULED

# *United States v. Heppner.*

There is a third exposure area most boards have not yet registered. It landed in federal court eight weeks ago, and it changes how every employee's prompts must now be treated.

CASE FILE

## *United States v. Heppner*, No. 25-cr-00503-JSR (S.D.N.Y.)

FIRST IMPRESSION NATIONWIDE

Judge Jed S. Rakoff · oral ruling Feb 10, 2026; written opinion Feb 17, 2026.

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. **The documents Heppner created with Claude were not protected by attorney-client privilege, and were turned over to federal prosecutors.** The court itself described it as a question of first impression nationwide.

### THE COURT'S REASONING · THREE PARTS

**01 An AI tool is not a lawyer.**

Claude has no law license, owes no duty of loyalty, and cannot form an attorney-client relationship. The court wrote that this fact alone "disposes of Heppner's claim of privilege."

**02 Public AI is not confidential.**

Privilege requires that a communication was intended to be, and actually was, kept confidential. Anthropic's own terms permit data collection and disclosure to governmental authorities.

**03 You cannot fix it after the fact.**

Privilege must exist at the moment of the communication. Routing something through your attorney *after* typing it into a public AI does not restore a privilege that was never there.

### THE DISTINCTION THAT GOVERNS YOUR LEGAL EXPOSURE

**PRIVATE AI** YOU CONTROL ALL THREE

The hardware it runs on, the model itself, and the data that flows through it. **What goes in stays in.**

**PUBLIC AI** YOU CONTROL NONE

ChatGPT, Claude.ai, Google Gemini, Microsoft Copilot consumer. Input goes to the vendor's servers, runs on their infrastructure, and is governed by terms that typically permit data collection, storage, and 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.

A FOURTH IMPLICATION THE PRESS LARGELY MISSED

# 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.

> **One careless prompt can potentially strip protection from an entire category of legal advice.**

*Print this. Read it to your senior team.*

WORK PRODUCT, ALSO REJECTED

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.**

IMMEDIATE IMPLICATIONS

- Any employee who uses public AI to analyze a legal question, draft a response to a regulator, or research a compliance issue may be creating a **discoverable record** opponents can obtain.
- This applies to civil litigation, workplace investigations, regulatory inquiries, and internal 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.**
- Sending the AI output to your general counsel afterward does not fix it.

08 THE RIGHT QUESTION FOR BOARDS NOW

# Not *whether* to use AI for legal work. Which tier, under what terms, with what oversight.

The privilege analysis differs materially across three tiers. Boards need to know which tier each of their workloads sits in — and which tier their employees *think* they're using.

TIER 1 — PUBLIC AI

## Heppner applies directly.

EXAMPLES

ChatGPT, Claude.ai, Google Gemini, and free or standard consumer versions of Microsoft Copilot.

LEGAL POSTURE

No reasonable expectation of confidentiality. Vendor terms permit data collection, training use, and disclosure to regulatory authorities.

**Privilege is destroyed at the moment of input.**

TIER 2 — ENTERPRISE AI

## Verified contractual confidentiality.

ALL THREE REQUIRED

- → Enterprise license at the appropriate tier.
- → Signed Data Protection Addendum.
- → Confirmed zero-training configuration — *verified*, not assumed.

RESIDUAL GAP

Indemnification compensates Microsoft's liability, not yours. Shared compute can still expose data through misconfigurations or side-channel risks no contract anticipates.

TIER 3 — PRIVATE AI

## The strongest position available today.

A model the company owns, running on hardware the company controls, with no external data flow. **No third party. No vendor terms. No disclosure risk.**

THE LIMITATION THAT SURVIVES

An AI is still not a lawyer, no matter who owns the hardware. Conversations with private AI, standing alone, are not privileged.

**Work product protection *can* apply — if the tool is deployed at counsel's direction, used for specific legal strategy work, and the relationship is documented.**

THE CONFIGURATION GAP

Most companies that "have Copilot" are running it under terms that do *not* provide the protections they assume. The license alone is not enough. **If your general counsel cannot point to documentation confirming each of the three Tier 2 conditions, you are still in Tier 1 for legal purposes** — regardless of what you paid for the license.

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. At most companies today, that means nearly everyone.

09 THE EXPOSURE SURFACE

# AI is bigger than cybersecurity — in shape, not just scale.

In the cyber era, the board's oversight obligation was primarily reactive: is the company managing breach risk and incident response? AI introduces a fundamentally different 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.

### 01 Algorithmic bias claims

Particularly in employment, lending, and consumer-facing applications subject to anti-discrimination law. State-level frameworks (California, Colorado) already create explicit disclosure obligations on top of federal exposure.

### 02 Model drift & error accumulation

An AI system's outputs degrade over time in ways not immediately visible to management. The risk is accelerating: Google's *Nested Learning* paradigm (NeurIPS 2025) is designed to make continuous in-model learning the default.

**Governance frameworks built around periodic audits will not keep pace.**

### 03 Vendor dependency & data leakage

Institutional knowledge gets embedded in a third-party model the company does not control. The IAC explicitly flagged third-party AI dependencies as a disclosure area regulators expect boards to assess.

### 04 Autonomous decision liability

When AI acts without a human decision-maker in the loop, accountability becomes legally ambiguous. The agent sent the email. The agent approved the credit. *Who answers for it?*

PATTERN

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

WHAT'S NEXT

Which brings us to the six questions every board should be able to answer — clearly, specifically, and not in boilerplate.

10

A DIAGNOSTIC

# Six questions every board should be able to answer.

Not in boilerplate. With specifics — named committees, documented cadences, real evidence. If your company cannot answer them clearly, **that gap is itself a governance finding**.

- **What AI systems does this company operate that could be considered material to investors, operations, or risk exposure?**
  **Translation:** *Have we inventoried them? Has any of them been disclosed? Is the inventory current?*
- **Has the board formally designated a committee responsible for AI oversight? Is that designation documented and reflected in the committee charter?**
  **Translation:** *If the answer is no, the IAC's "disclose or explain" standard now requires you to say why not.*
- **What information does the board receive about AI risk — how often, in what format, from whom? Is that process documented?**
  **Translation:** *A board that cannot show a written information-flow is a board that cannot show oversight.*
- **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?**
  **Translation:** *Disclosure accuracy is the second liability path. Stale language drafted in 2023 is not a defense.*
- **Has legal counsel reviewed our AI-related risk factor disclosures in light of the IAC recommendations and SEC staff guidance from 2024–2025?**
  **Translation:** *Not as a one-time exercise — as part of a documented annual review cycle.*
- **Have we assessed our third-party AI dependencies — which vendor models we rely on, what data we are providing to them, and whether that creates training data or competitive risk?**
  **Translation:** *The IAC and the Heppner ruling both pointed at third-party AI. Where are your Tier-2 verifications?*

### A NOTE ON THE PACE OF CHANGE

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

The appropriate board posture is not to wait for formal rules. It is to build governance structures and disclosure practices now that will hold up under *whatever* framework emerges.

### THE LINE THAT MATTERS

**The companies that wait are betting that regulators, plaintiffs' attorneys, and institutional investors will not ask hard questions before the rules are fully formed.**

**Based on the past two years, that is not a safe bet.**

ABOUT THIS PAPER

# Authored by Iterate's executive team.

This paper was prepared by Iterate's executive team, with research drawn from Claude (Anthropic) and cross-referenced by Iterate's Private AI *Generate* platform. It was prepared for educational purposes and reflects publicly available legal and regulatory information as of April 2026.

**This document is not legal advice and should not be relied upon as such.** Laws and regulations in this area are evolving rapidly. Consult qualified legal counsel before making decisions based on the information contained here.

AUTHORS

**Magnus Tagtstrom** — CORP VP

Worked on GDPR-compliant products in Europe for Circle K.

**Dave Jenkins** — VP, RESEARCH

Ran OEM partnerships for Red Hat in Asia and EMEA.

**Jon Nordmark** — CEO

Served on Colorado's governor- and legislature-appointed AI Task Force during the development of the first broad-sweeping US state AI law.

REFERENCES

01 SEC Division of Corporation Finance, *Statement on AI Disclosure Priorities* (Eric Gerding), June 24, 2024.

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

03 SEC Investor Advisory Committee, *"Disclosure of Artificial Intelligence's Impact on Operations,"* Dec 4, 2025.

04 SEC Division of Examinations, *2025 Examination Priorities* (Oct 2024); *2026 Examination Priorities* (Nov 2025).

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

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

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

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

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

10 Google Research, *Nested Learning* — NeurIPS 2025.

**A final note on scope.** The SEC framework described here is one layer of a rapidly expanding legal landscape. Iterate's broader Board AI Governance session 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. Directors and executive teams should expect legal exposure to expand well beyond SEC disclosure as AI law matures across jurisdictions.

SERIES

**Handout 02**

Board AI Governance

HEADQUARTERS

**San Jose, CA**

& Denver, CO

ON THE WEB

**iterate.ai**

hello@iterate.ai

EDITION

**April 2026**