



The Center for
AI Oversight

CENTER PAPER NO. 3 · JUNE 2026

The Committee Fallacy

The anatomy of the AI oversight program: what the committee must govern, and the difference between having a committee and governing through one.

By **Brian J. Allen**, Executive Director

Each section moves in three tiers: a **board line**, a **why it matters** note, and a shaded **detail** block carrying the authority.

EXECUTIVE SUMMARY

The committee is the governing body, not the governance

KEY TAKEAWAYS

- 1 The most common board response to AI risk, assign it to a committee, is the precise error Delaware has adjudicated twice. The committee is the governing body; it is not the governance.
- 2 *Clovis* held a board liable despite a committee, because it did not monitor. *Boeing* held a board liable because no charter named the risk. The charter that is silent fails; the charter that speaks but is not followed fails identically.
- 3 Charters are adding the word AI faster than programs are being built beneath them. That moves institutions from Boeing's exposure to Clovis's, a change of venue, not progress.
- 4 What the committee must govern is the program: five pillars, fourteen domains, answering what is governed and who is accountable, with how left to the operational layers.
- 5 The objection that the program slows AI down is backwards. Ambiguous obligation is the brake; clear obligation is what lets AI move. The program is the governance.

This paper names the error, documents how quickly institutions are committing it, specifies the anatomy of the program the committee must actually govern, and answers the speed objection directly. The layering is precise because the precision is what makes governance defensible: the committee governs the program, the program guides the AI activity, and the record proves both.

PART ONE

The fallacy, named and documented

1.1 Why committee designation feels like governance

Assigning a risk to a committee produces a documentable act and signals attention, but neither is governance. Governance is a continuing activity; the charter amendment is its precondition, not its performance.

Why it matters. The distinction would be pedantic if Delaware had not built two cases on it. Clovis is the controlling instruction: a board with a committee overseeing its lead clinical trial was liable because it knew of reporting problems and did nothing. Boeing is the threshold half: no charter named the risk, and silence read as oversight absence. Only the charter that is followed, in documented committee activity, protects.

CASES

In re Clovis Oncology, Inc. Derivative Litig., C.A. No. 2017-0222-JRS (Del. Ch. Oct. 1, 2019); *In re The Boeing Co. Derivative Litig.*, C.A. No. 2019-0907-MTZ (Del. Ch. Sept. 7, 2021). Together they bracket the fallacy from both sides: the silent charter and the unfollowed charter fail alike.

1.2 The gap, measured

Institutions are naming and assigning AI risk far faster than they are building the program beneath it. The charter amendment without a program moves the institution from the first prong of the doctrine to the second.

Why it matters. The first prong, no system at all, is becoming unavailable to plaintiffs. The second prong, a system not monitored, is where the AI cases will be fought, and the disclosure data shows institutions walking straight into it.

DATA

31.6% of the S&P 500 disclosed some board oversight of AI as of the 2024 proxy season, up more than 84% in a single year; explicit disclosure of full board or committee oversight stood at 11% (ISS-Corporate, March 2025), and a March 2026 review found board-level AI oversight disclosed at 8% across the Russell 3000 and S&P 500. Meanwhile 72% of the S&P 500 disclosed AI as a material risk in 2025 filings, 83% by April 2026 (The Conference Board / ESGAUGE).

The charter amendment moves the institution from Boeing's exposure to Clovis's. The first prong asks whether the structure exists. The second asks whether it governs. Structure is the easy prong.

PART TWO

The anatomy of the program

2.1 What the committee governs

The committee governs the program: a defined, documented system establishing what must be governed and who is accountable, leaving how to management, counsel, and operators. Crossing into the how dissolves the committee's position above the activity.

Why it matters. The Center's published architecture gives the anatomy a name and structure: five pillars, fourteen domains. The pillars and domains are public; the measurement layer beneath them is the work of the Center's assessment instruments. The anatomy alone is sufficient here, because the anatomy is what the case law describes from the outside.

2.2 The five pillars and fourteen domains

Pillar I. Agile Governance (the constitution)The foundational management system: the program's mandate, authority, and adaptability. D1 AI Governance Program & Policy Framework (the mandate); D2 AI Governance Structure, Oversight & Resources (accountability and authority, where the committee charter and officer ownership live); D3 Governance Program Assurance & Continuous Learning (validating that the program operates and evolves).

Pillar II. Risk-Informed System (the guardrails)The repeatable methodology for identifying, assessing, and communicating AI risk against board-approved tolerance. D4 AI Risk Methodology, Scope & Tolerance (the yardstick and the program's reach); D5 Risk Intelligence & Threat Landscape (forward-looking intelligence that keeps the board's information current).

Pillar III. AI Trust & Assurance (the evidence)The standards AI systems must meet, on the principle that trust is earned through evidence, not assertion. D6 AI Model Risk & Agent Lifecycle Oversight; D7 AI Data Governance Oversight; D8 AI Transparency, Explainability & Human Oversight; D9 AI Security & Resilience Assurance.

Pillar IV. Risk-Based Strategy & Execution (strategic alignment)The pillar linking governance to business objectives so oversight improves decisions rather than impeding them. D10 Risk-Informed Strategy, Resources & Organizational Readiness; D11 AI Value Realization & Operational Resilience Oversight; D12 Third-Party AI & Supply Chain Governance (where the Buyer position's obligations enter the program).

Pillar V. Risk Escalation & Disclosure (the voice)The disciplined flow of material information to leadership. D13 AI Risk Escalation & Disclosure Protocols (critical risks bypass silos and reach decision-makers, on defined triggers and clocks); D14 Validation of Escalation & Governance Effectiveness (testing the emergency brakes and auditing the decision process).

2.3 The anatomy against the case law

The architecture answers Boeing element by element. The committee with charter responsibility is Domain 2; the reporting cadence and protocols are Domains 2 and 13; the independent channel is Domain 13's design requirement; the documented red-flag response is Domain 13 in execution and Domain 14 in verification.

Why it matters. The case law describes the program's absence; the program is the case law's specification, run in reverse and in advance. The overarching question the program teaches its committee to ask is not whether a specific model is safe, but whether the program is designed and operating effectively so the institution can make informed decisions on risk and reward. That is the Informed Decision Standard in interrogative form.

MAPPING

Marchand's duty to name and inventory the mission critical risk is the work of Domains 4, 6, and 12; Clovis's monitoring failure is what Domains 3 and 14 exist to prevent, where internal audit's mandate under the IIA Global Internal Audit Standards meets the program.

QUESTIONS FOR THE BOARD

- Has our committee been given charter responsibility for AI in the charter's text, and can its minutes show it governing, not merely meeting?
- Which of the fourteen domains can we evidence today, and which exist only on an org chart?

PART THREE

The objection: will the program slow us down

The objection that fourteen domains means bureaucracy is backwards. Ambiguous obligation is the brake on AI adoption; clear obligation is what lets AI move.

Why it matters. Where no one knows who can approve a deployment, what standard it must meet, or what happens when it misbehaves, every decision escalates to everyone and velocity collapses. The deployment that meets the Pillar III standards, inside the Domain 4 appetite, with the Domain 13 triggers attached, can be approved by the accountable owner without convening the institution. The program is the reason the institution can say yes quickly and defend the yes later.

LEGAL PRICING

Texas makes documented NIST AI RMF alignment a statutory defense; Connecticut credits documented anti-bias testing in mitigation; the Marriott dismissal shows the documented program defeating a Caremark claim after a catastrophic outcome; and examiners in half the states now expect a written insurer AI program. The institutions that built the program bought speed and protection with the same expenditure.

CONCLUSION

The program is the governance

The committee that governs the program is doing what Delaware, the examiners, and the investors are all separately asking to see. The committee that merely exists is the fallacy with a charter.

The anatomy is not mysterious and never was: a mandate, a yardstick, standards, strategic alignment, and a voice; five pillars, fourteen domains; what is governed and who is accountable, with how left to the layers built for it. The doctrine has been explicit about which one it tests, and the disclosure data shows the test approaching at scale. The program is the governance.

Sources and Further Reading

Cases: *In re Clovis Oncology, Inc. Derivative Litig.*, C.A. No. 2017-0222-JRS (Del. Ch. Oct. 1, 2019); *In re The Boeing Co. Derivative Litig.*, C.A. No. 2019-0907-MTZ (Del. Ch. Sept. 7, 2021); *Marchand v. Barnhill*, 212 A.3d 805 (Del. 2019); *In re McDonald's Corp. Stockholder Derivative Litig.*, 289 A.3d 343 (Del. Ch. 2023); *Firemen's Ret. Sys. of St. Louis v. Sorenson*, C.A. No. 2019-0965-LWW (Del. Ch. Oct. 5, 2021).

Other authority and data: Tex. Bus. & Com. Code chs. 551-552; Connecticut Pub. Act 26-15; NAIC Model Bulletin (Dec. 4, 2023); IIA Global Internal Audit Standards; ISS-Corporate (Mar. 2025); ISS Governance QualityScore (Mar. 2026); The Conference Board / ESGAUGE (Oct. 2025; updated Apr. 2026). The AI Oversight Program architecture is the Center's published framework; doctrinal treatment of the cases appears in full in *The Caremark AI Liability Roadmap*.