WHAT I LEARNED COMPARING WHAT COMPANIES SAY IN PUBLIC TO WHAT THEY DOCUMENT FOR REGULATORS

An independent review of how four leaders talk about AI — to the public and investors, and in what they file with the SEC

About the Analysis

To better understand how organizations are approaching AI governance and Board oversight, I conducted an independent review of publicly available governance materials from four large, publicly traded medtech and life science companies.

I began with publicly available investor presentations, where these organizations discussed AI strategy, innovation priorities, and expected business outcomes. I also read the AI principles they publish on their own websites.

Then I went to what they file with regulators. Each year these same companies file a 10-K and a proxy statement with the SEC, describing the same business and the same AI for a different audience, under a different standard, where inaccuracy carries legal consequences. I read each company's 10-K and proxy, and I compared all of it against what they had said in public.

My review covered sixteen primary-source documents.

Research Question

As AI becomes a strategic priority for public companies, do their governance, oversight, and accountability disclosures keep pace with the AI narrative they present to investors and the public?

While the organizations reviewed were all leaders in their sector, the analysis was not intended to identify best practices. It was intended to identify patterns.

The Observation That Surprised Me

Across four large public companies, no common model for Board oversight of AI emerged. Sometimes AI governance belonged in the full Board; in other cases, oversight was distributed across different Board committees. There was no common theme in AI governance management structures or reporting lines. The implication is not that one model is right and another wrong — it is that a company may need to be far more deliberate about where AI governance resides, and why, based on its strategy, AI risk, and operating model.

Sources Reviewed

[Small graphic — 4 companies × 4 document types = 16 primary sources]

	Investor Conference	Public AI Principles	10-K	Proxy
Company A	✓	✓	✓	✓
Company B	✓	✓	✓	✓
Company C	✓	✓	✓	✓
Company D	✓	✓	✓	✓

All findings drawn from SEC filings on sec.gov, company investor relations sites, and corporate websites.

Four Observations

One — AI governance does not live in one place. Across these four companies, AI oversight is housed in four structurally different ways: distributed across several board committees, assigned to full-board annual review, run through a management-level center of excellence, or anchored in audit and finance. There is no industry standard yet, and no obviously right answer. The question for any board is whether its structure was chosen deliberately or inherited by default.

Two — The proxy, not the 10-K, is where AI governance actually shows up. Both are filed with the SEC, but they do different jobs. The 10-K tells you what a company is worried about. The proxy tells you how the board is governing it. Three of the four companies use the proxy as the primary venue for naming AI as a board-level matter. A reader who looks at only one of those documents sees only half the picture.

Three — Data is named everywhere as the foundation of AI risk, but the harder question goes unanswered. Every 10-K names data quality or data fitness as a determinant of AI risk. Yet few disclosures provide meaningful visibility into whether the underlying data is reliable, traceable, explainable, and fit for the purpose for which the AI is being used. This matters because data can satisfy traditional quality controls and still be unsuitable for an AI application. The question this raises is whether the controls Boards should ask about for AI are operationally different from the ones they have historically asked about for cybersecurity and privacy.

Four — Dependence on a small number of AI infrastructure providers is rising quickly, and no single filing is built to make it visible. Foundation-model partnerships and accelerated-compute relationships are being announced at a fast pace, and the disclosures acknowledge that a company building on a third party's AI model may have limited rights to the underlying intellectual property, which limits its ability to independently verify explainability and reliability. This concentration becomes obvious only when you read across companies at once. A board governing AI in this environment needs a third-party assessment framework that is operationally different from a traditional vendor-risk review.

Underneath all four is one connective theme. These documents are written by different teams, for different audiences, under different standards. When they tell the same story, governance, oversight, and strategy reinforce one another. When they drift apart, stakeholders are left to reconcile competing narratives about how AI is being governed.

Questions for Management and Boards

These four questions follow naturally from the comparison. They are not a scorecard. They are the questions I would want answered if this were my company.

1. **Where does AI live on your board, and is the placement deliberate?** If AI risk has been inherited by your cybersecurity committee through the SEC's Item 1C structure, was that an active decision or a default?
2. **Read together, do your 10-K, proxy, AI principles, and latest investor materials tell the same story about AI strategy and AI risk?** Who is responsible for reconciling them, and how often does that actually happen?
3. **How does management demonstrate that the data supporting AI initiatives is reliable, fit for purpose, and governed appropriately?**
4. **For each named third-party AI relationship, does your vendor-risk framework address underlying-IP access, explainability, and concentration risk?** Is your contract language calibrated for vendors whose technology is woven into regulated products?

To Consider Next

These four questions lead to a harder one. If a board asked management to *prove* its AI governance was working, what would a credible answer look like? That question matters enough that I took it up on its own, in a companion review of what the standards bodies, the board-governance institutions, and institutional investors say boards should be asking management to measure, and how management should be prepared to respond with evidence.