

The Board Director's Playbook: Responsible Director Use of Artificial Intelligence

By **Shellye Archambeau**, Operating Partner, Head of the Board Impact Practice at The ExCo Group, and Lauren L. Valiente

Published October 26, 2026



Artificial intelligence (AI) is moving from the enterprise into the boardroom — helping directors research, draft, summarize, and test assumptions. That convenience carries a governance cost: a single prompt can expose material nonpublic information and privileged communications, and a fluent answer that is materially wrong can influence board decision-making. This playbook focuses on directors' own use of AI, not management's enterprise AI program, and provides directors with a practical framework for use that is informed, secure, reviewable, and candid.

I. Introduction

AI is now part of ordinary governance work: preparing for meetings, organizing diligence, testing assumptions, comparing scenarios, drafting questions, and reviewing public information. It can increase efficiency, but it also adds fiduciary, confidentiality, disclosure, privilege, cybersecurity, and recordkeeping risk because the prompt itself may contain information the company must protect.

The governing principle is simple: AI may assist a director's preparation, but it cannot replace the director's judgment, confidentiality obligations, or duty to become adequately informed. Delaware law places the corporation's business and affairs under the board's direction, 8 Del. C. § 141(a), and protects good-faith reliance on corporate records and appropriate information, opinions, reports, or statements when the statutory conditions in 8 Del. C. § 141(e) are met. The objective is not to ban useful tools but to treat them as controlled boardroom instruments — not confidential advisers, independent experts, substitutes for counsel, or invisible board members.

II. A Case Study Illustration

The following hypothetical illustrates the most common risks posed by a director's use of AI:

Sam is an independent director of ABC, Inc., a Delaware corporation listed on the NYSE. ABC is evaluating a proposed acquisition of a competitor that is not yet public. The board is reviewing detailed information about the proposed acquisition, including the competitor's identity, valuation and financing assumptions, management projections, negotiation status, diligence findings, and potential timing. The information is confidential and includes material nonpublic information about ABC and the proposed acquisition.

While riding in an Uber to the board meeting, Sam wants to use AI to test his thinking before the discussion. The company-approved AI platform is unavailable on his phone, so he uses a personal ChatGPT account instead. He enters detailed information about ABC and the proposed acquisition — including the competitor's identity, valuation, financing, projections, diligence, negotiations, and timing — and asks ChatGPT to identify the principal risks, assess whether the proposed acquisition appears justified, and recommend how he should approach the discussion.

ChatGPT returns a confident but materially flawed analysis. It overlooks an important diligence concern, accepts a key transaction assumption without support, and gives Sam an overly reassuring conclusion about the proposed acquisition. Sam then quotes that conclusion during the board discussion without disclosing that it was AI-generated.

During the board meeting, the board is informed of an anonymous complaint submitted through ABC's whistleblower hotline alleging a material misstatement in ABC's last earnings release. The board's discussion includes speculation about the facts and who may be responsible. During the meeting, Sam utilizes an AI program to listen to the discussion and generate a summary after the meeting. After the meeting, Sam also uses ABC's approved AI platform to generate an analysis of the relative strengths and weaknesses of ABC's legal position based on the board's discussion.

Sam may believe he acted efficiently and in ABC's interests. His conduct nevertheless creates three distinct, overlapping risks. First, he disclosed confidential information about the proposed acquisition to a personal ChatGPT account outside ABC's controls. Second, he allowed a materially flawed AI output to influence the board process without verifying its assumptions or reliability. Third, he created exposure by entering and generating sensitive information about legal risk into an AI Platform that will be discoverable in litigation and by regulators, including the Securities and Exchange Commission (SEC) and the Department of Justice (DOJ). The first risk implicates confidentiality, cybersecurity, and disclosure controls; the second implicates an informed and candid board process; and the third implicates privilege and discovery. The sections that follow address those risks in that order and then translate the lessons into policy and oversight rules.

III. Risks Created by Sam's Conduct

1. Confidentiality and Disclosure Controls

Sam's disclosure of detailed information about ABC's proposed acquisition into a personal ChatGPT account was not merely an informal use of a device. It placed ABC's confidential information outside the approved platform and may have exposed it to retention, human access, model training, subprocessors, or a security incident. That creates confidentiality, contractual, cybersecurity, and disclosure-controls risk. ABC should assess the incident under its disclosure controls and procedures, including Exchange Act Rules 13a-15(e) and 15d-15(e), and under its incident-response process.

Sam's use of a Personal ChatGPT Account does not place the conduct outside ABC's governance perimeter. ABC's code of conduct, confidentiality procedures, disclosure controls, incident-response plan, and records policies should apply, regardless of whether Sam used a personal account, personal device, or consumer service. The resulting records should be handled under the same governance framework as other company communications.

2. The Board Process Must Remain Demonstrably Informed

Sam's materially flawed output creates a separate board-process problem. The issue is not that AI assisted his preparation; it is that he quoted its reassuring conclusion without disclosure and did not test or verify the assumptions before allowing it to shape the discussion. Delaware law permits a director to rely in good faith on corporate records and information from appropriate sources when the statutory conditions of Section 141(e) are satisfied; an undisclosed, unverified consumer-AI output should not be treated as that kind of expert input. ABC may therefore have difficulty showing that Sam's approach to the proposed acquisition was based on a reasonable understanding of its risks. AI Use Policies should require director disclosure to the board when AI materially influences a director's questions, recommendation, risk assessment, or vote.

3. Privilege, Work Product, and Discovery

Sam's later query in the Approved Platform regarding the whistleblower complaint may serve a legitimate business purpose, but it is not privileged merely because he utilized a platform that was approved to request a "legal analysis." A prompt to an AI system is not a communication with counsel. Attorney-client privilege only applies to communications between a client and their attorney for the purpose of rendering legal advice, and work product protection only applies to material produced by attorneys, or at their direction, in anticipation of litigation. Neither of those protections apply when directors ask AI for a legal assessment of information learned in a board meeting.

Sam's use of an AI program to summarize the meeting discussion is similarly problematic. Most AI programs that summarize meetings also record and retain a complete audio recording and written transcript of the meeting. In most cases, those materials will be fully discoverable to regulators and to litigants.

AI prompts, uploaded files, and generated outputs are fast becoming common subjects of discovery requests in litigation, collection plans in internal investigations, and DOJ and SEC subpoenas. Directors should assume that prompts and outputs may be discoverable regardless of whether they use the company's Approved AI Platform and even when the use served a legitimate purpose. Once a duty to preserve this data arises, intentional destruction of any data can create spoliation or obstruction risk, including sanctions and, in appropriate circumstances, criminal exposure.

Directors should apply the same discipline to AI activity that they apply to email, text messages, board portals, and handwritten notes: use approved systems, preserve material communications, avoid personal accounts, do not destroy records after a hold, and assume that informal communications may be reviewed outside the context in which they were created. Practical safeguards include: (i) do not paste any sensitive legal information including legal advice, investigation material, or whistleblower allegations into any AI platform; and (ii) notify the general counsel or records custodian immediately if a prompt may be responsive to a hold, investigation, subpoena, Section 220 request, or regulatory inquiry.^{ns}

IV. Why Directors Need an AI-Use Policy

An AI-Use Policy fills a gap that enterprise controls cannot close. A company may have an AI governance program for employees and vendors, yet directors may use personal devices, personal email, consumer applications, or tools provided by another board. Sam's use of a Personal ChatGPT Account because the Approved AI Platform was unavailable on his phone illustrates why directors need a clear rule for what to do when the approved tool is inconvenient or inaccessible.

An AI-Use Policy also gives the director a practical process for satisfying the duty of care. It helps distinguish low-risk assistance from high-risk use involving proposed acquisitions, litigation, investigations, whistleblower reports, financial reporting, or other sensitive board matters. This risk-based approach is consistent with the Department of Justice's Evaluation of Corporate Compliance Programs, which asks whether compliance programs are adequately resourced, capable of monitoring emerging technology, and

updated as risks change. National Institute of Standards and Technology (NIST) AI Risk Management Framework 1.0 can provide a voluntary framework for governing, mapping, measuring, and managing AI risk, but it does not replace legal advice or board judgment.

The policy should be specific enough to answer the questions that matter. At a minimum, it should tell the director:

- Permitted tools, accounts, and use cases.
- Information that must never be entered.
- When human review, disclosure, escalation, or counsel involvement is required.
- What AI records must be retained and how to respond to a hold or investigation.

V. What the Policy Should Include

Section III identifies the risks; the policy must convert them into operating rules. A director's personal AI-use policy should apply whenever a director uses AI in connection with any company business including the board, a committee, a proposed transaction, an investigation, a shareholder matter, or preparation for a board discussion.

1. **Scope and definitions.** Define "AI," "generative AI," "Approved Platform," "board information," "material nonpublic information," and "AI reliance." Clarify that the policy covers prompts, uploaded files, outputs, browser extensions, plug-ins, mobile applications, transcription tools, meeting assistants, and AI features embedded in ordinary software.
2. **Approved tools and accounts.** Identify the Approved Platform and limit a director's permitted use of AI for company business to the Approved Platform. Prohibit consumer tools and applications that route prompts through unknown third parties. If the Approved Platform is unavailable on a device, defer the task, use a secure approved alternative, or contact the designated support or legal contact; do not substitute a personal AI account.
3. **Prohibited inputs.** Unless expressly authorized on a platform reviewed and approved by counsel and security personnel, prohibit entry of any nonpublic business information including board books, draft minutes, financial results or forecasts, litigation strategy, legal advice, work product, personal data, trade secrets, customer information, source code, credentials, security architecture, and information subject to a confidentiality agreement
4. **Confidentiality, disclosure, and data security.** Prohibit any use that discloses or exposes material nonpublic information, confidential board materials, privileged communications, or personal data to an unauthorized recipient. State expressly that approval of the Approved Platform does not eliminate disclosure, preservation, discoverability, or production obligations.
5. **Attorney-client privilege and work product.** Require consultation with the general counsel before using AI to generate any type of legal assessment or uploading any investigation material, litigation strategy, whistleblower allegations, or material prepared at counsel's direction. Use only a platform with terms

and controls approved for the relevant use case and assume that any information provided to an AI platform, even company approved AI platforms, is not privileged and will be discoverable.

6. **Disclosure of material AI reliance.** Require disclosure when AI materially influences a director's recommendation, risk assessment, or vote. Distinguish material reliance from low-risk uses, such as summarizing public information. Treat quoting an AI conclusion in a board discussion without identifying its source as a material process concern when it influenced the director's position.
7. **Prompt and output retention.** Adopt a retention and deletion schedule that was reviewed and approved by counsel. Specify where records are stored, who may access them, how long they are retained, and how a director must respond to a litigation hold, investigation notice, subpoena, Section 220 demand, or regulatory inquiry.
8. **Personal accounts, mobile access, and public settings.** Prohibit use of personal AI accounts for confidential board work. If the Approved Platform is unavailable on a mobile device, wait, use a secure approved alternative, or seek assistance rather than use a personal AI account.
9. **Prohibit the use of AI to transcribe or summarize board meetings.** Prohibit all uses of AI programs to transcribe or summarize board or committee meetings. Require the immediate removal of any AI participant from board or committee meetings.
10. **Vendor terms and no-training commitments.** Before approving a platform, confirm the vendor's data-use, retention, deletion, model-training, subprocessors, location, security, audit, confidentiality, and breach-notification terms. Require a contractual commitment that board data and prompts will not be used to train or improve a generally available model without express authorization, and address ownership of inputs, outputs, logs, metadata, and derived data. Confirm how records can be preserved and produced if requested by regulators or litigants.
11. **Escalation.** Require prompt escalation to the general counsel, chief information security officer, or designated board contact if a director enters prohibited information, discovers unauthorized access, or learns that a vendor retained or disclosed data.
12. **Training, certification, enforcement, and insurance.** Require annual training and a written director acknowledgment or certification. State that violations may trigger remedial training, suspension of platform access, committee review, or other action consistent with the company's governance documents. The board should also confirm that its Directors and Officers (D&O) insurance, indemnification arrangements, and notice procedures are understood; coverage should not be assumed for conduct involving unauthorized disclosure, spoliation, or intentional misconduct.

VI. Ongoing Oversight

The board or designated committee should review the AI-Use Policy periodically — at least annually, and sooner when AI tools, director use cases, or material risks change. The review should test whether the policy remains clear, practical, and current; identify provisions that need updating as technology, security

practices, and regulatory developments evolve; and confirm that any vendor or platform review supports, rather than substitutes for, policy oversight.

VII. Practical Rules for Directors' Use of AI


These rules are designed for the moment when a director is about to open a chatbot, upload a document, or rely on an AI-generated answer. They translate the policy into simple stop-and-check actions:

1. Use only the Approved AI Platform through a company-managed or board-approved account.
2. Never enter material nonpublic information into an unapproved tool; use the Approved Platform only when the policy permits its specific use.
3. Never use an AI program to transcribe and summarize board or committee meetings.
4. Immediately remove any AI participants from meetings.
5. Do not enter nonpublic company information into personal AI accounts.
6. Do not use public Wi-Fi, rideshare settings, shared devices, or voice assistants for sensitive prompts.
7. Use AI to organize, compare, or generate questions — not to decide whether a risk is material or how you should vote.
8. Disclose material AI reliance to the board before discussion and before a vote.
9. Do not delegate director judgment to AI; independently verify every material analysis, assumption, and source before relying on it.
10. Ask the general counsel before using AI for legal advice or uploading legal materials to an AI platform.
11. Assume everything you type into a prompt and every document you upload can be discovered by regulators or in litigation.
12. Follow the company's retention policy with respect to the retention of any AI inputs, including prompts and uploaded documents, and AI generated outputs.
13. If a hold or investigation begins, preserve the full AI record and follow counsel's instructions; do not self-help or selectively preserve only the favorable output.
14. Complete the required certification annually and whenever the policy, Approved Platform, or board role changes materially.

This article was originally published by [Foley & Lardner, LLP](#) in partnership with [Shellye Archambeau](#), Exco Operating Partner and Head of the Board Impact Practice.

Read the original [here](#).

About the authors



Shellye Archambeau

Operating Partner, Head of the Board Impact Practice



Lauren L. Valiente

The ExCo Group is an exclusive community of experienced CEOs, board directors, and global business leaders who bring wisdom, executive judgment, and intellectual rigor to coaching and leadership development. Our mentors and coaches are former CEOs and global business leaders who work exclusively with us to help executives, teams, and boards outperform.

More leadership insights from The ExCo Group: excoleadership.com/articles/