Society for Corporate Governance

Structured Chatham House Rule Discussion on Legal Privilege and AI

May 20, 2026

Key Takeaways

The Society for Corporate Governance convened a structured discussion on **Legal Privilege and AI** on May 20, 2026. The program, conducted under the Chatham House Rule, brought together 124 Society members to discuss, among other topics, when and how attorney-client privilege may apply in an AI context; how AI platform terms and data practices affect confidentiality; the role of counsel direction in preserving privilege and work product protection; and risks of discoverability of AI prompts and outputs.

David Gordon and Takayuki Ono, Partners at Sidley Austin, served as featured speakers and used recent court decisions as a framework for examining these issues in practice. Drawing on emerging case law, they: (i) explored how different courts have analyzed privilege and work product claims involving AI prompts, outputs, and related materials—reaching different outcomes based on specific facts and circumstances; (ii) identified factors that may influence future judicial outcomes; and (iii) discussed practical considerations for organizations seeking to manage legal, governance, and discovery risks associated with AI use, including tool selection, internal policies, and training. Society members in attendance contributed additional perspectives to the discussion.

Below are key takeaways from the discussion. *These takeaways are not, and should not be construed as, legal advice.*

- **Organizations should assume that traditional privilege and work product principles will generally apply to the use of AI, including AI-generated outputs and related work product.** As illustrated by *United States v. Heppner* and *Warner v. Gilbarco*, courts evaluating privilege and work product issues involving AI have largely relied on existing legal doctrines rather than creating AI-specific legal standards. As a result, longstanding legal principles such as confidentiality, legal purpose, attorney involvement, anticipation of litigation, and waiver remain central to the analysis.

- **AI-related privilege and discoverability issues are a rapidly evolving area of law, and outcomes depend heavily on the specific facts and circumstances.** Recent decisions, including in *Heppner* and *Warner*, provide useful guidance but demonstrate that courts may reach different conclusions based on how AI was used, who used it, the purpose of the interaction, and the safeguards in place to protect confidentiality. As courts continue to apply existing legal doctrines to rapidly developing technologies, organizations should adopt policies and practices that are designed to withstand evolving legal standards across jurisdictions and evaluate AI use in its specific context rather than relying on broad assumptions about the technology.

- **The confidentiality terms associated with an AI platform may significantly affect whether privilege protections apply.** In *Heppner*, the court found that communications with Anthropic's consumer Claude AI platform were not protected by attorney-client privilege, relying in part on the platform's terms permitting disclosure of user information to third parties, including law enforcement. The decision underscores the importance of understanding how AI providers collect, retain, use, and disclose prompts and outputs, as well as the potential legal significance of the distinction between *consumer* AI tools (i.e., publicly available AI platforms designed for general use and typically governed by standard terms of service) and *enterprise* AI tools (i.e., AI services provided under business agreements that generally include enhanced confidentiality, security, administrative controls, and data-use restrictions).

- **Attorney involvement may be an important factor in determining whether AI-assisted work receives privilege protection.** In *Heppner*, the court rejected privilege claims for litigation-related materials developed through an AI tool without attorney involvement, but suggested that a different result might be possible where the tool is used at the direction of counsel and functions in a manner similar to a trained professional acting as an attorney's agent. The decision highlights the potential risks of relying on AI tools as a substitute for legal counsel when addressing legal issues or developing litigation-related analyses.

- **The use of AI by in-house counsel for legal purposes may not be fundamentally different from the use of other legal research and productivity tools, provided appropriate safeguards are in place.** Privilege and work product protections are likely to apply where in-house counsel uses an enterprise AI platform subject to confidentiality protections, employs the tool in connection with legal work, and restricts access to legal personnel involved in the matter. Maintaining separate AI workflows for legal activities (vis-à-vis business activities) and clearly identifying the legal purpose of AI-assisted work—including, where appropriate, within AI prompts themselves—may further support privilege claims.

- **The scope of work product protection for AI-assisted litigation materials likely will depend on context.** In *Heppner*, the court rejected work product protection for materials prepared without counsel's involvement, even though they were created in anticipation of litigation. By contrast, in *Warner*, the court found that a pro se plaintiff's ChatGPT-assisted litigation strategy materials were protected work product, reasoning that the doctrine protects a party's thought processes and mental impressions and that use of an AI tool did not necessarily constitute disclosure to an adversary. These differing approaches suggest that work product protection for AI-generated materials may heavily depend on the applicable jurisdiction, the role of counsel, and the circumstances under which the AI tool was used.

- **AI prompts and outputs may become discoverable evidence in litigation.** As illustrated by *Fortis Advisors v. Krafton*—in which a CEO's ChatGPT interactions became part of the evidentiary record and were cited by the court in evaluating the company's conduct, decision-making, and intent—AI prompts and outputs may be scrutinized in court alongside other electronically stored information. Organizations should therefore assume that AI interactions may be subject to preservation, review, and production obligations and could be used to assess intent, knowledge, and reasoning behind corporate actions.

- **Use of AI can create unique discoverability and litigation risks by preserving early-stage and unguarded thinking that might not otherwise be documented.** Employees often turn to AI tools at the outset of a problem, project, investigation, or decision-making process, before conducting deeper analysis or consulting others. As a result, prompts and outputs may capture preliminary impressions, incomplete analyses, candid questions, or untested assumptions that were never intended to represent final conclusions, recommendations, or corporate positions. Because users often engage with AI in a conversational manner similar to a search engine, these records may be more direct, speculative, and less carefully considered than emails or other business communications, creating heightened litigation and governance risks.

- **Reliance on AI outputs without clearly defined evaluation criteria, meaningful human oversight, and appropriate documentation can create significant legal and governance risks.** As illustrated by *American Council of Learned Societies v. National Endowment for the Humanities*, the use of vague prompts, the absence of meaningful human review, and insufficient documentation regarding AI's role in the process may make it difficult to validate or defend resulting decisions. In addition, when AI is used in connection with legal matters, the absence of clear prompts or documentation regarding the purpose and use of the AI-assisted analysis may make it more difficult to establish or defend claims of privilege. Organizations may therefore benefit from establishing clear parameters for AI-assisted analyses, documenting how AI contributes to decision-making, and ensuring that human reviewers independently evaluate AI-generated recommendations before action is taken.

- **Organizations should exercise particular caution when using AI recording, transcription, and summarization tools in board, legal, executive, and other sensitive meetings.** Unlike traditional meeting minutes, AI-generated transcripts and summaries may create detailed records of discussions that were never intended to be preserved verbatim. Such tools may introduce discoverability, accuracy, consent, privacy, confidentiality, and other legal risks, particularly where sensitive strategic, legal, or governance matters are discussed. While retention controls may help mitigate some of these risks, they should be aligned with legal hold and preservation obligations. Organizations should therefore carefully evaluate when these tools are appropriate and establishing clear policies governing their use.

- **Effective AI policies do more than prohibit the use of unauthorized tools—they should help employees identify when legal involvement is required.** Organizations should make employees aware of the risks associated with using unauthorized or consumer AI tools, emphasize that AI is not a substitute for legal counsel, especially in legally sensitive matters, and train employees to recognize when legal guidance should be sought before using AI tools or relying on AI-generated outputs. Without clear guidance, employees may inadvertently use AI tools in connection with litigation, investigations, regulatory inquiries, board matters, transaction-related activities, compensation decisions, or other legally sensitive matters, potentially creating privilege, confidentiality, and discoverability risks. AI policies and training programs will therefore be most effective when they clearly identify situations requiring the involvement of legal counsel and provide employees with practical, role-specific guidance on the appropriate use of AI in those contexts.

- **Efforts to govern directors' use of AI may be more effective when driven by board leadership and peer norms rather than formal restrictions alone.** Because directors often operate with significant independence, organizations may have limited ability to control how individual directors use AI tools. Aligning key board leaders on AI-related risks, reinforcing expectations through education, and framing AI use within familiar governance concepts—such as the risks associated with inappropriate email communications, note-taking practices, or unauthorized recordings—may be more effective than relying solely on restrictions. Organizations should emphasize that certain uses of AI can create risks not only for the individual user but for the board and company as a whole.

Program Materials

- Sidley Slide Deck
- Generative AI and Privilege: Practical Lessons from Two Early Decisions and What Comes Next
- When "The Devil Made Me Do It" Is Not a Defense: Lessons in AI Governance and Organizational Oversight from an SDNY Decision
- AI in the Boardroom: Governance Implications and Fiduciary Considerations
- Society Slide Deck

Access additional AI-related resources—including sample policies, template guidelines and frameworks, basic resources, and benchmarking surveys—on our Artificial Intelligence page.