

AI-Related Issues in Securities Cases: Privilege Pitfalls and ‘AI Washing’ Claims

WRITTEN BY

Jay A. Dubow | Erica Hall Dressler | Milica Krnjaja

This article was originally published on [The Legal Intelligencer](#) and is republished here with permission as it originally appeared on May 26, 2026.

Artificial intelligence (AI) tools have quickly become integral to users across the spectrum, from large public companies to individual consumers. Public companies now routinely reference “AI-driven” strategies on earnings calls, and boards are being briefed on AI risks alongside cybersecurity and ESG. The legal industry is no exception. Law firms and their clients increasingly rely on AI to draft disclosures, summarize documents, and prepare for litigation. Given this widespread use and popularity, AI-related issues have unsurprisingly begun to feature prominently in recent securities cases. These issues range from allegations of “AI washing” and overstated AI capabilities to disputes over whether clients’ use of generative AI tools can be shielded by the attorney-client privilege.

DISCOVERY RULINGS

On the privilege front, Judge Jed Rakoff’s recent opinion in *United States v. Heppner* shows how quickly AI use can become a discovery problem in a securities case. Bradley Heppner, a former executive of publicly traded GWG Holdings, Inc., was indicted in October 2025 for securities fraud, wire fraud, conspiracy, making false statements to auditors, and falsifying corporate records, based on an alleged scheme to defraud GWG investors by causing GWG to enter undisclosed, self-serving transactions with entities he controlled. After his arrest, the FBI executed a search warrant at his home and seized documents and electronic devices. Among the seized materials were approximately 31 documents memorializing his exchanges with a publicly available generative AI platform.

According to defense counsel, these “AI documents” were created in 2025 after Heppner received a grand jury subpoena and understood he was a target of the investigation. He allegedly used the AI platform on his own initiative to prepare reports outlining potential defense strategies, possible arguments about the facts and law, and anticipated government theories. Counsel later asserted privilege over these documents, arguing that Heppner had input into the AI information he learned from his lawyers, created the documents for the purpose of speaking with counsel and obtaining legal advice, and then shared the contents with counsel. The government moved for a ruling that the AI documents were not protected by the attorney-client privilege or the work-product doctrine, and Judge Rakoff granted that motion.

In analyzing privilege, the court held that the AI documents were not communications “between a client and his or her attorney” because the AI platform is not an attorney and no attorney-client relationship exists with such a tool. The court also concluded that the documents were not confidential. The platform’s privacy policy expressly

allowed collection of user inputs and outputs, use of that data to train the model, and disclosure to various third parties, including regulators, in connection with disputes or litigation—undermining any reasonable expectation of confidentiality. In addition, the communications were not made to obtain legal advice from the AI platform, which itself disclaims providing legal advice and directs users to consult an attorney.

On work product, the court stressed that the doctrine protects materials prepared by or at the behest of counsel that reflect counsel’s mental processes. The AI documents were generated by Heppner “on his own volition,” not at counsel’s direction, and while counsel later said the documents “affected” their thinking, they did not “reflect” counsel’s strategy when created. Under U.S. Court of Appeals for the Second Circuit precedent, that was not enough to bring them within work-product protection. The court thus concluded that Heppner’s AI “defense reports” were discoverable.

Heppner, however, arose in the criminal context, and some civil courts have expressly distinguished its holding when extending work-product protection to AI-assisted materials. For instance, in *Morgan v. V2X*, a civil employment case in the U.S. District Court for the District of Colorado, Magistrate Judge Dominguez Braswell held that the work-product protections in Federal Rule of Civil Procedure 26(b)(3) can apply to a pro se litigant’s use of AI tools in preparing for litigation. The court reasoned that Rule 26(b)(3) protects materials prepared “by or for” a party, not just attorneys, and that AI-assisted notes and analyses can reflect the party’s mental impressions and litigation strategy. The court nevertheless required the plaintiff to disclose the identity of the AI tool used—finding that naming the tool did not itself reveal protected mental impressions—and amended the protective order to prohibit uploading “Confidential” information into mainstream AI platforms absent strict contractual safeguards.

The court in *Morgan* relied on the recent decision of the U.S. District Court for the Eastern District of Michigan in *Warner v. Gilbarco*, which also involved a pro se plaintiff’s use of AI tools. In *Warner*, the defendants sought production of “all documents and information” concerning the plaintiff’s use of third-party AI tools and argued that any work-product protection had been waived by using a public generative-AI platform. The court rejected that argument, holding that a pro se litigant may invoke Rule 26(b)(3) to shield AI-assisted materials that reflect her mental impressions and litigation strategy, and that work-product protection is not automatically waived by transmitting those materials to an AI vendor because such platforms are “tools, not persons,” and disclosure to them is not the same as disclosure to an adversary or in circumstances likely to place the materials in an adversary’s hands.

Notably, both *Morgan* and *Warner* extended work-product protection to AI-assisted materials in civil cases brought by pro se plaintiffs, whereas *Heppner*—arising in the criminal context and involving a represented defendant—held that similar AI “defense reports” were discoverable. This divergence shows that the law in this area is still developing, and it is not yet clear how courts will apply these principles in securities class actions and other complex civil cases involving represented corporate clients that increasingly rely on AI tools in their day-to-day legal and disclosure functions. Consequently, persons using AI tools in litigation matters should consider the risk that such use may be discoverable to the opposition.

AI-RELATED MISREPRESENTATION CLAIMS

In addition to discovery rulings relating to AI, we have continued to see AI-related allegations serve as the primary theory in securities-fraud and related enforcement cases. A recent example is the criminal prosecution of iLearningEngines Inc.’s former CEO and CFO in the U.S. District Court for the Eastern District of New York.

iLearningEngines, a Maryland-based company, marketed itself as providing an AI training platform for schools and health care entities. In April 2026, prosecutors unsealed a 10-count indictment alleging that the executives “exploited investor excitement over the AI boom” by fabricating the company’s customer relationships, using sham contracts to create the illusion of real AI clients and to inflate reported revenue and earnings. The government further alleges that they “round-tripped” funds between lenders, purported customers, and the company to make revenue appear legitimate, and that the company later filed for bankruptcy after a short-seller report claimed it had overstated its Indian revenue by 99%. The case remains pending, and no trial date has yet been set.

Investors filed a parallel civil suit in Maryland federal court against iLearningEngines and its executives based on similar allegations, but that case was dismissed in March 2026 for failure to state a claim.

Another recent case involves a publicly traded financial technology platform that uses data and AI to support consumer lending. In April 2026, a shareholder filed a securities class action in the U.S. District Court for the Northern District of California under Sections 10(b) and 20(a) of the Exchange Act and Rule 10b-5, alleging that the company and certain executives misled investors about the performance of its latest AI-driven approval model, “Model 22.” The company launched Model 22 in May 2025 and publicly touted its accuracy, stating that it would increase loan approvals and revenue, and, according to the complaint, raised its annual revenue guidance in both May and August 2025 on the strength of those claims.

On Nov. 4, 2025, however, the company announced third-quarter results that missed its guidance and lowered its outlook for the fourth quarter and full year. On the same day, the company attributed the shortfall to Model 22, explaining that the model had “overreacted” to negative macroeconomic signals, that the company had “knowingly” calibrated it to be more conservative on credit risk, and that its “over-responsiveness” would continue to depress approvals and revenue. The stock price declined by roughly 10%. The complaint alleges that the company failed to disclose that Model 22’s risk-separation processes frequently overreacted to macroeconomic signals, that its purported accuracy and approval-rate benefits were overstated, and that the model’s conservatism was already undermining the company’s revenue guidance.

AI practices themselves, not just revenue and model-performance claims, are also beginning to generate follow-on D&O litigation. For instance, on April 24, 2026, a shareholder filed a derivative action in the U.S. District Court for the Northern District of California against the board and senior executives of Adobe Inc., asserting breach of fiduciary duty and waste, as well as claims under Section 14(a) and Section 10(b) based on alleged misstatements in proxy materials and public disclosures about Adobe’s AI practices. According to the complaint, Adobe’s SlimLM small language model tools relied on datasets such as “Books3” and “Common Crawl,” which allegedly contain hundreds of thousands of copyrighted works obtained without authorization. The plaintiff contends that Adobe’s directors and officers “followed an ‘ask forgiveness, not approval’ model” by choosing not to use “clean” datasets, and that this strategy led to multiple copyright lawsuits, a share price decline of more than 25% after the first IP suits were filed, and the March 2026 departure of CEO and Chairman Shantanu Narayan, allegedly linked to Adobe’s “failed AI strategy.” The case remains in its early stages and is still pending.

CONCLUSION

The discovery landscape around AI use is evolving rapidly, and the cases discussed here are early snapshots rather than settled law. Courts are still working through how existing privilege and work-product doctrines apply

when AI is embedded in both corporate operations and legal practice, leaving users of AI with multiple unanswered questions about best practices. At the same time, plaintiffs attorneys, regulators, and prosecutors are continuing to apply familiar securities fraud and disclosure frameworks to AI-centric fact patterns, from alleged “AI washing” and overstated model performance to AI-driven IP and governance failures. Thus, companies and their advisers should assume that both their use of AI tools and their AI-related disclosures will continue to be judged in real time for the foreseeable future, and should build policies and practices with that reality in mind, ensuring they are flexible enough to adapt as this area of law continues to develop.

Reprinted with permission from the May 26, 2026 edition of *The Legal Intelligence* © 2026 ALM Global Properties, LLC. All rights reserved. Further duplication without permission is prohibited, contact 877-256-2472 or asset-and-logo-licensing@alm.com

RELATED INDUSTRIES + PRACTICES

- [Artificial Intelligence](#)
- [Business Litigation](#)
- [White Collar Litigation + Investigations](#)