View on our website.

GIBSON DUNN

White Collar Defense & Investigations Update February 20, 2026

AI Privilege Waivers: SDNY Rules Against Privilege Protection for Consumer AI Outputs

Gibson Dunn's Artificial Intelligence, Privacy, Cybersecurity, & Data Innovation, and White Collar Defense & Investigations Practice Groups are available to advise on best practices, contractual strategy, training initiatives, and risk management considerations at the intersection of AI, privilege, and related legal exposures.

Overview

On February 10, 2026, Judge Jed S. Rakoff of the Southern District of New York, in an oral ruling from the bench, held that materials generated through a consumer AI tool at the prompting of a criminal defendant were not protected by either the attorney-client privilege or the work product doctrine.[1] In its subsequent February 17 written opinion explaining that decision, the Court characterized the issue as a "nationwide" matter of first impression—namely, whether communications with a publicly available AI platform during a pending criminal investigation are protected by privilege or work product.

The Court concluded that attorney-client privilege protection was unavailable because: (1) the AI tool was not a lawyer and could not establish an attorney-client relationship; (2) there was no expectation of confidentiality because the platform's privacy policy disclosed that user inputs and platform outputs could be used for model training and disclosed to third parties; and (3) the defendant did not communicate with the AI tool for the purpose of obtaining legal advice. The Court also concluded that the work product doctrine was unavailable because the defendant generated the AI materials independently, rather than at the direction of counsel.[2] In reaching

its decision, the Court made clear that AI's "novelty" does not mean its use "is not subject to longstanding legal principles."[3]

Judge Rakoff's ruling highlights the litigation risks associated with employee or individual use of consumer AI tools without careful attention to governing terms of service and data practices. Users of publicly accessible or "open" AI platforms may assume that privilege attaches to their inputs and the resulting AI outputs; but a platform's privacy policy permitting data collection, model training, or disclosure to third parties may defeat any reasonable expectation of confidentiality or privilege. Importantly, however, the Court did not announce a rule uniquely targeting AI technologies; rather, it applied traditional attorney-client privilege and work product legal principles to the conduct at issue.

Background

As recounted in the parties' filings: On October 28, 2025, a federal grand jury indicted Bradley Heppner on securities and wire fraud charges.[4] In the months preceding the indictment, Mr. Heppner received grand jury subpoenas and had been informed that he was a target of the Government's investigation. Although he retained counsel in anticipation of indictment, Mr. Heppner—acting independently and not at counsel's direction—submitted prompts to an AI platform's conversational model to organize and synthesize information he believed relevant to his defense.[5] Certain inputs reflected information he had learned through discussions with counsel. He used the tool to generate written materials intended to consolidate his thoughts and facilitate future communications with counsel.[6]

On November 4, 2025, when Mr. Heppner was arrested, the FBI executed a search warrant at his residence and seized electronic devices containing materials generated through his interactions with the AI tool.[7] Defense counsel subsequently identified 31 documents as reflecting Mr. Heppner's AI-generated analyses and asserted those documents were protected from disclosure by the attorney-client privilege and work product doctrine.[8] On February 6, 2026, the Government moved for a determination that the 31 AI-generated documents were not subject to protection from disclosure. At the hearing on February 10, 2026, Judge Rakoff ruled from the bench that there was no basis for the defense's privilege and work product claims.

Basis for Decision

At first glance, Judge Rakoff's ruling may prompt concern regarding the continued reliability of privilege and work product protections in the context of AI-assisted work. A close reading of the transcript and written opinion,[9] however, suggests a narrower holding. The Court's determination was fact-specific and grounded in conventional privilege and work product legal principles, applied to a factual scenario shaped by the distinct contractual and technological features of the specific consumer AI platform in use, as well as the defendant's unilateral decision to utilize the AI platform without direction from counsel.

Attorney-Client Privilege

With respect to attorney-client privilege, the Court concluded that Mr. Heppner could not satisfy several required elements.

First, the Court held that the Mr. Heppner's AI generated documents were not communications between Mr. Heppner and his counsel. Although the Court acknowledged ongoing debate about whether a user's AI inputs, rather than being communications, are more "akin to the use of other Internet-based software, such as cloud-based word processing applications," the Court explained that use of such applications are likewise "not intrinsically privileged."[10] Rather, the inquiry turns on the presence of a "trusting human relationship," such as in the context of attorney-client privilege—a relationship with "a licensed professional."[11]

Second, the Court concluded that Mr. Heppner "had no 'reasonable expectation of confidentiality in his communications'" with the AI platform.[12] That determination rested not merely on the fact that the defendant had communicated with a third-party AI platform, but on the tool's governing privacy policy, which disclosed that the platform collects users' "inputs" and "outputs," uses that data to "train" its model, and reserves the right to disclose such data to multiple categories of "third parties," including the government. Under those circumstances, the Court concluded that the communications were not confidential.

Third, the Court found that Mr. Heppner did not communicate with the AI tool for the purpose of obtaining legal advice.[13] Mr. Heppner's counsel argued that his client communicated with the AI tool for the "express purpose of talking to counsel." But the Court pointed out that the tool's terms expressly disclaimed that it is not a lawyer and cannot provide formal legal advice, and when asked directly, the AI tool responded that it could not provide legal advice.[14] The Court acknowledged that had counsel directed Mr. Heppner to use the AI tool to aid in subsequent attorney-client communications, the result may have been different.[15] In doing so, the Court was careful to warn that even if Mr. Heppner intended to share the AI-generated documents with counsel, communications that are not privileged at the time they are made do not "acquire protection merely because they were transferred" to counsel.[16] Nor could they "somehow alchemically change[] into privileged" material by later being shared with counsel.[17]

Work Product Doctrine

The Court separately addressed the work product doctrine. Although the defense advanced a work product argument based on Mr. Heppner's inclusion of information obtained from counsel, the Court emphasized that the AI-generated materials were not prepared at counsel's direction.

Defense counsel confirmed that the AI-generated documents were "prepared by the defendant on his own volition," and the Court therefore concluded that Mr. Heppner was not acting as counsel's agent when he communicated with the AI platform.[18]

The Court further distinguished between materials that merely "affect" litigation strategy and those that "reflect" counsel's mental impressions at the time of their creation. While the AI-generated documents created by Mr. Heppner may have influenced counsel's thinking going forward, they did not embody or memorialize counsel's strategic analysis when they were generated.[19] For that additional reason, they did not qualify for work product protection.

A Traditional Application of Privilege in the AI Era

Importantly, Judge Rakoff's ruling does not represent a categorical rejection of privilege and work product in the AI context. Rather, it reflects the application of longstanding privilege waiver and work product principles to a modern technological intermediary—an application made more complex by evolving assumptions about how AI platforms handle user-submitted information, and subject to continuing discourse on whether AI outputs constitute entirely "new" information or are new versions of the initial user inputs.

Takeaways

The practical complication is that many users may underestimate the extent to which certain AI systems are architected around the ingestion, review, and reuse of user inputs to improve model performance. In functional terms, engaging with some consumer AI platforms may resemble confiding in a third party who expressly reserves the right to disseminate or repurpose what is shared and created. From a privilege perspective, that structural reality cannot be ignored.

Absent negotiated contractual protections, protective privacy and confidentiality settings, and other technical guardrails, consumer AI tools may function—from a privilege perspective—as third parties that retain, review, and leverage user-submitted information to train and improve their models; such facts may undermine claims of confidentiality for purposes of privilege and work product protection.

The *Heppner* ruling underscores that baseline AI literacy and disciplined legal risk management are essential. Organizations and individuals can take concrete steps to mitigate privilege and confidentiality risks while still capturing the operational benefits AI tools offer:

- **Invest in education**. Stakeholders should understand that many AI business models are premised on provider access to, and potential reuse of, user prompts and platform outputs. Misapprehensions about how these systems function can lead to inadvertent waiver or confidentiality lapses.

- **Scrutinize contractual terms**. Service agreements should be reviewed carefully to assess data access, retention, and secondary-use provisions. Where AI tools will be used for sensitive or legally significant matters, organizations should consider enterprise deployments or "closed" environments that restrict provider access and prohibit model training on submitted data. Individual and non-enterprise users, by contrast, should assume that standard consumer terms typically do not provide such protections.

- **Closely review publicly available policies**. AI providers' publicly stated terms of service, privacy policies, or other public statements should be examined for disclosures that could constitute waivers of confidentiality. Even where such statements are not part of agreements directly entered with users, contradictory statements could undermine future privilege claims.

- **Assess and use AI settings**. AI providers' access to user inputs for use in training AI models is often enabled by default. However, most major AI platforms allow users to opt out of having their inputs (such as uploads, prompts, and conversations) used for training models. There may be limitations to opt-outs via AI platform settings so the underlying

terms should be fully assessed. A more fulsome evaluation of terms of service and other contractual terms is preferable, but opt-out settings are a good starting point to ensure better protections.

- **Use AI deliberately and with governance controls**. AI can serve as a force multiplier when deployed strategically. Unstructured use without proper AI governance and policy controls—particularly in connection with communications relating to legal advice or anticipated litigation—can create unnecessary privilege exposure. Clear internal policies and appropriate involvement of legal counsel are critical to preserving protections. Use of AI tools during or in anticipation of litigation should be done solely at the direction of counsel to provide more durable privilege protections.

- **Use proper descriptions in litigation**. If AI outputs are recorded on privilege logs or otherwise in dispute in discovery, counsel should clearly indicate that the AI tools were used at counsel's direction, and that the outputs reflect attorney mental impressions and were created under circumstances supporting a reasonable expectation of confidentiality.

The central takeaway from Judge Rakoff's ruling is not that AI adoption is incompatible with privilege and work product protections, but that unexamined use of AI tools can create avoidable legal risk. Thoughtful evaluation, contractual diligence, and structured deployment can substantially mitigate those risks.

[1] February 10, 2026 minute entry, *United States v. Heppner*, No. 25-cr-00503-JSR (S.D.N.Y.) (*Heppner*) (granting Dkt. 22 motion for a ruling that documents the defendant generated through an AI tool are not privileged); Transcript of February 10, 2026 Pretrial Conference at 6, *Heppner*.

[2] Dkt. 27 at 5–7, *Heppner*.

[3] *Id.* at 12.

[4] Dkt. 3, *Heppner*.

[5] Dkt. 22 (motion for a ruling that documents the defendant generated through an AI tool are not privileged) at 7–8, *Heppner*.

[6] Dkt. 23-4 (notes of 1/21/26 call) at 7–8, *Heppner*; Dkt. 23-5 (notes of 2/2/26 call), *id.*

[7] Dkt. 22 at 7–8, *Heppner*; Press Release, United States Attorney's Office Southern District of New York, Former CEO And Board Chairman Charged With Fraud Scheme Directed At Public Company (Nov. 4, 2025), https://www.justice.gov/usao-sdny/pr/former-ceo-and-board-chairman-charged-fraud-scheme-directed-public-company.

[8] Dkt. 23-2 (defense privilege log) at 7–8, *Heppner*.

[9] Dkt. 27, *Heppner*.

[10] *Id.* at 5.

[11] *Id.* at 6–8.

[12] Dkt. 27 at 7, *Heppner* (quoting *United States v. Mejia*, 655 F.3d 126, 132 (2d Cir. 2011)).

[13] *Id.* at 6.

[14] *Id.* at 7–8.

[15] *Id.* at 7.

[16] *Id.* at 8.

[17] *Id*.

[18] *Id.* at 9–10.

[19] *Id.*

The following Gibson Dunn lawyers prepared this update: Eric Vandevelde, Diana Feinstein, M. Theodore Takougang, Justin Lin, and Meredith Spoto.

Gibson Dunn's lawyers are available to assist in addressing any questions you may have regarding these issues. Please contact the Gibson Dunn lawyer with whom you usually work, the authors, or any leader or member of the firm's White Collar Defense & Investigations, Artificial Intelligence, or Privacy, Cybersecurity & Data Innovation practice groups:

Eric D. Vandevelde – Los Angeles (+1 213.229.7186, evandevelde@gibsondunn.com)

Diana M. Feinstein – Los Angeles (+1 213.229.7351, dfeinstein@gibsondunn.com)