Paul, Weiss

Client Memorandum

February 20, 2026

S.D.N.Y. Court Considers Whether AI-Generated Documents Are Subject to Privilege Protections

On February 10, 2026, Judge Jed S. Rakoff of the U.S. District Court for the Southern District of New York ruled that 31 documents that a defendant, Bradley Heppner, generated using a third-party AI tool (the "AI-Generated Documents") were not protected by either the attorney-client privilege or the work product doctrine. Judge Rakoff's ruling in the case, *U.S.* v. *Heppner*, appears to be the first in which a court determined that interactions with a publicly accessible AI tool based on prompts that contain privileged information are not themselves privileged.

At the pre-trial conference and in his opinion, Judge Rakoff rejected any theory of attorney-client privilege, reasoning, inter alia, that the defendant had disclosed information provided to him by his counsel to a public AI tool in which there was a diminished expectation of confidentiality.[1] The court further declined to treat the AI-Generated Documents as work product after determining that the defendant had not prepared the material at the direction of counsel.[2]

It remains to be seen how broadly the ruling will be applied. In his decision, Judge Rakoff explained that the defendant's documents were not confidential because: (1) Heppner "communicated with a third-party AI platform"; (2) the AI tool's written privacy policy, to which users like Heppner consent, "provides that [the company that owns the tool] collects data on both users' 'inputs' and the [tool's] 'outputs,'"; (3) the collected data is used to train the tool; and (4) the company "reserves the right to disclose such data to a host of 'third parties,' including 'government regulatory authorities.'"[3]

These findings suggest a highly fact-specific ruling. While the traditional contours of privilege as applied to this case should remain consistent going forward, Judge Rakoff's findings regarding the privacy policy may be of limited applicability. Companies, and potential litigants, should pay particularly close attention to the terms of any AI tool they wish to use, specifically whether the AI company uses data for training. That being said, beyond AI tools, it is fairly standard for organizations to have terms in their privacy policies to allow for disclosure to government regulatory authorities and other third parties, so companies and their counsel should keep an eye on how courts handle this finding going forward.

[1] Transcript of Record at 2:25–3:15, *U.S.* v. *Heppner*, 1:25-cr-00503-JSR (S.D.N.Y. Feb. 10, 2026); *See* Memorandum, *U.S.* v. *Heppner*, 1:25-cr-00503-JSR (S.D.N.Y. Feb. 17, 2026) (No. 27) [hereinafter "Opinion"], at 6.

[2] Tr. at 5:2–6:21; Opinion at 9.

[3] Opinion at 6.

Background: The Attorney-Client Privilege and the Work Product Doctrine

As a general matter, the attorney-client privilege "protects communications (1) between a client and his or her attorney (2) that are intended to be, and in fact were, kept confidential (3) for the purpose of obtaining or providing legal advice."[4] The party that is invoking the privilege bears the burden of demonstrating that it applies to the communications at issue.[5]

The work product doctrine protects materials that are prepared by counsel, or at counsel's direction, in anticipation of litigation or for trial.[6]

The AI-Generated Documents and the Government's Motion

On November 4, 2025, defendant Bradley Heppner was arrested on fraud charges. At that time, law enforcement seized a host of electronic devices from his residence.[7] Defense counsel notified the government that, before his arrest, the defendant had used Anthropic's Claude, an AI tool, to run a series of queries related to the government's investigation of his conduct.[8] Acting without any suggestion or direction from counsel, Heppner used Claude to "prepare[] reports that outlined defense strategy, that outlined what he might argue with respect to the facts and the law that [h]e anticipated that the government might be charging."[9] His counsel asserted privilege over the AI-Generated Documents, arguing: (1) Heppner's inputs into Claude included information that he learned from counsel, (2) Heppner "had created the [AI-Generated] Documents for the purpose of speaking with counsel to obtain legal advice," and (3) Heppner shared the documents with his counsel.[10] The parties subsequently entered a privilege protocol stipulation to segregate the AI-Generated Documents from the prosecution team pending resolution of the privilege question.[11]

With trial less than two months away, the government moved for the court to rule that the AI-Generated Documents were not protected by the attorney-client privilege or the work product doctrine. The government contended that the AI-Generated Documents "fail every element" of the attorney-client privilege,[12] arguing:

- The AI-Generated Documents were not privileged communications between the defendant and his counsel since Claude "is obviously not an attorney" and owes no duties of loyalty and confidentiality to its end users.[13]
- The AI-Generated Documents were not created for the purpose of obtaining legal advice because, if asked about giving such advice, Claude responds that it cannot do so and that users should consult qualified attorneys.[14]
- The documents were not confidential because Anthropic's privacy policy states that it collects data on user activity to train its models and that it may also disclose such data to regulatory bodies and third parties. As a result, according to the government, the defendant had a diminished expectation of privacy. [15]
- The defendant's later transmission of the AI-Generated Documents to his counsel could not retroactively create privilege.[16] While the government acknowledged that any advice that counsel offered in response to the AI-Generated Documents

[4] *United States* v. *Mejia*, 655 F.3d 126, 132 (2d Cir. 2011).

[5] *In re Grand Jury Subpoenas Dated Mar. 19, 2002 and Aug. 2, 2002*, 318 F.3d 379, 384 (2d Cir. 2003).

[6] *Id.* at 383.

[7] Opinion at 3.

[8] *Id.*

[9] *Id.*

[10] *Id.* at 3–4.

[11] *Id.* at 4.

[12] Motion for a Ruling that Documents the Defendant Generated Through an AI Tool Are not Privileged at 7, *U.S.* v. *Heppner*, 1:25-cr-00503-JSR (S.D.N.Y. Feb. 6, 2026) (No. 22) [hereinafter "Motion"],at 1.

[13] *Id.* at 7.

[14] *Id.* at 1, 8.

[15] *Id.* at 9.

[16] *Id.* at 9–10.

would be presumptively privileged, the government argued that the protection does not extend to the information in the underlying documents.[17]

The government then turned to its argument that the AI-Generated Documents are not protected by the work product doctrine.[18] The government raised defense counsel's apparent concession that the defendant had used Claude of his own accord and not at the direction of counsel.[19] According to the government, the defendant's decision to transmit his independent research to counsel did not transform the AI-Generated Documents into protectable work product created with an eye toward litigation.[20]

Defense counsel did not file a response to the government's motion ahead of the pre-trial conference.

The Pre-Trial Conference and Judge Rakoff's Decision

Judge Rakoff concluded that the AI-Generated Documents were not privileged communications for three reasons: (1) Claude is not an attorney, (2) the AI-Generated Documents were not confidential, and (3) Heppner did not use Claude to obtain legal advice.

First, Judge Rakoff determined that the AI-Generated Documents did not constitute "communications between Heppner and his counsel." Claude "is not an attorney," and "Heppner d[id] not, and indeed could not, maintain that Claude is an attorney." In fact, Judge Rakoff continued, all "recognized privileges" require a "trusting human relationship," and no relationship "with a licensed professional who owes fiduciary duties and is subject to discipline . . . exists, or could exist, between an AI user and a platform such as Claude."[21]

Second, the AI-Generated Documents were not confidential because Heppner communicated with a "third-party AI platform," and he consented to Anthropic's privacy policy. According to the privacy policy, Anthropic collects data on user activity to train its models and "reserves the right to disclose such [users'] data to" third parties, such as regulatory authorities. As a result, Anthropic's policy put users like Heppner "on notice that Anthropic, even in the absence of a subpoena . . . may 'disclose personal data to third parties.'" [22] In addition, Judge Rakoff stated that Heppner had "no reasonable expectation of confidentiality in his communications" with Claude because, according to another recent Southern District of New York case, "AI users do not have substantial privacy interests in their 'conversations with [another publicly accessible AI platform], which users voluntarily disclosed' to the platform and which the platform 'retains in the normal course of its business.'"[23] The AI-Generated Documents were "not like confidential notes that a client prepares with the intent of sharing them with an attorney because Heppner first shared the equivalent of his notes with a third-party, Claude."[24]

On this finding, the ruling seems to rely on the government's assertion in its briefing that the public AI tool trains on "prompts" and "outputs" and may disclose the relevant data.[25] However, the policy cited states "[w]e will not use your Inputs or Outputs to train our models, unless: (1) your conversations are flagged for Trust & Safety review. . . , or (2) you've explicitly reported the materials to us . . . , or (3) you've otherwise explicitly opted in to the use of your Inputs and Outputs for training purposes."[26] Additionally, the policy's language relating to disclosure to government authorities and third parties with respect to legal obligations is similar to language commonly found in privacy policies, including those for many law firms and technology companies. [27] As such, the impact of the decision here may be limited to the specific facts at issue.

[17] *Id*.

[18] *Id.* at 11.

[19] *Id.*

[20] *Id.*

[21] Opinion at 5.

[22] *Id.* at 6.

[23] *Id.* at 6–7.

[24] *Id.* at 7.

[25] Motion at 9.

[26] Anthropic Privacy Policy (Feb. 19, 2025), https://www.anthropic.com/legal/archive/a2eecf43-807a-4a53-89dd-04c44c351138

[27] *Id.*

Finally, the AI-Generated Documents were not privileged because "Heppner did not communicate with Claude for the purpose of obtaining legal advice." Although Judge Rakoff acknowledged that this issue was a "closer call," Heppner did not use Claude "at the suggestion or direction of counsel," and the Documents did not become privileged simply because they were shared with counsel. Further undercutting defense counsel's argument was Claude's explicit disclaimer in response to a Government inquiry about whether it could give legal advice: "I'm not a lawyer and can't provide formal legal advice or recommendations." Claude then recommended that a user consult with a qualified attorney.[28] Judge Rakoff noted, however, that were the situation slightly different—if counsel had directed Heppner to use the AI tool—"Claude might arguably be said to have functioned in a manner akin to a highly trained professional who may act as a lawyer's agent within the protection of the attorney-client privilege."[29]

Judge Rakoff also ruled for the government on the issue of whether the AI-Generated Documents constituted work product. Defense counsel argued that materials the defendant created in anticipation of his indictment that were intended for counsel—even if not made at their direction—should receive protection.[30] Defense counsel noted that the defendant began querying Claude after receiving a grand jury subpoena and after discussions with the government made clear that he was the target of an investigation.[31] Defense counsel therefore argued that the "purpose of [the defendant] preparing these reports was to share them with [counsel]" to discuss his defense.[32]

Judge Rakoff rejected this argument, stating that "[t]he core purpose of Work Product Doctrine is to protect the mental strategies of counsel in anticipation of litigation."[33] In his opinion, Judge Rakoff explained that even assuming that the AI-Generated Documents were "prepared in anticipation of litigation," they were not "prepared by or at the behest or counsel" and "nor did they reflect defense counsel's strategy," especially given the fact that Heppner used Claude on his own volition.[34]

Key Takeaways

- **Ordinary Privilege Principles Applied**. Judge Rakoff and all of the parties applied the existing, fundamental principles of the attorney-client privilege to AI use, including that the communication at issue must be kept confidential. Here, the court noted the privacy policy for the AI tool put users like the defendant on clear notice that their personal data may be disclosed to third parties, reiterating that "AI users do not have substantial privacy interests in their" conversations with publicly available AI tools when they share such data voluntarily and the platform retains the data in the normal course of its business.[35] The decision emphasizes the need for lawyers and non-lawyers alike to think carefully about what information they enter into, and receive from, AI tools if they expect to withhold those materials later.

- **The Nature of the AI Tool May Matter**. The exact AI tool at issue—including the tool's terms and design—may impact decisions involving AI use and privilege. While Judge Rakoff noted that the AI tool at issue disclosed prompts to the tool's developer, AI tools in which inputs and outputs are kept confidential could be treated differently. Users should examine an AI tool's terms, including provisions about whether the developer may access, train on, or share inputs and outputs. Enterprise-licensed AI tools will generally provide greater confidentiality than consumer-facing tools. In addition, AI tools that do not share any data with a third party may likewise provide stronger arguments that inputs and outputs are kept confidential. This may be the case for tools that operate (i) completely within a network environment that is private to a business or (ii) on an individual's device, such as with local or edge AI computing.

- **The Use Case for the AI Tool May Matter**. The use case of the AI tool may also affect privilege determinations. For the attorney-client privilege to apply, the communication must be between a client and their attorney. As Judge Rakoff noted, the use of "Internet-based software, such as cloud-based work processing applications" is not "intrinsically privileged."[36] However, the government acknowledged that there are certain exceptions to this requirement, which they claimed were not

[28] Opinion at 7–8.

[29] *Id.* at 7.

[30] Tr. at 3:17–25.

[31] *Id.* at 4:9–17.

[32] *Id.* at 4:21–23.

[33] *Id.* at 5:2–4.

[34] Opinion at 9–10.

[35] Opinion at 6–7.

[36] *Id.* at 5.

relevant to the defendant.[37] In particular, under *U.S.* v. *Kovel*, communications to a third party which are deemed essential to facilitate communication between the attorney and the client do not waive privilege.[38] Future actions will likely provide guidance on whether using an AI tool in support of obtaining legal advice, such as to analyze complex financial documents or translate foreign language materials, would be subject to the attorney-client privilege.

- **Taking Steps to Strengthen Work Product Claims**. In support of a claim for work product protection, when a non-lawyer assists with a litigation or prosecution at the direction of counsel and uses an AI tool, the non-lawyer should separately document that fact. Such records should clearly state that the party is querying an AI tool at counsel's direction. As noted above, non-lawyers (and lawyers) should also consider whether the nature of the particular AI tool being used could create the risk that "disclosure" of work product to the AI tool constitutes a waiver of work product protection.[39]

- **Issues Still Need to be Resolved**. Judge Rakoff's decision highlights that there are a host of other issues surrounding AI use in litigation. For example, while pro se litigants benefit from AI tools, their use may amplify some key risks, including the submission of hallucinated cases to the court. In addition, a potentially unintended consequence of using AI tools could be an increase in discovery requests relating to use of those tools by litigants in the context of a litigation. Further, Judge Rakoff's ruling leaves unanswered important questions about how other elements of the attorney-client privilege might apply to a client's use of an AI tool in connection with legal representation. There is still a threshold question whether a client can take the initiative to use an AI tool in support of obtaining legal advice or any action must be at counsel's direction. Other exceptions may be litigated in the AI context, too, such as where a third party is deemed to be the "functional equivalent" of a corporate employee.[40]

Conclusion

Judge Rakoff's ruling illustrates the risk of using publicly available technology tools, especially as law related to the use of AI is in its infancy and attempting to keep pace with rapidly developing technology. While here, Judge Rakoff declined to extend privilege protection to interactions with an AI tool by a party that had privileged information, different facts will likely lead to different results.

Businesses and individuals using enterprise-licensed AI tools or tools that do not share any data with a third party will likely have stronger arguments about confidentiality than those using consumer-facing AI products. Even then, confidentiality may further turn on the specific reason an individual is using an AI tool and whose privilege is at issue. An employee using an enterprise-licensed AI tool to analyze complex documents to discuss with in-house counsel may have an expectation of confidentiality supporting the *company's* privilege. If that same individual uses the enterprise-licensed AI tool, which his employer can monitor, to help his own attorney with a personal lawsuit, a court is more likely to find a lack of confidentiality defeating the *employee's* privilege.

Regardless, parties and counsel should be aware, especially as nascent case law develops, that courts may find that their use of AI tools can impact determinations of confidentiality and privilege.

[37] Motion at 7.

[38] 296 F.2d 918, 291–22 (2d Cir. 1961) (finding a client's communication to an accountant was privileged).

[39] Work product protection may be defeated if disclosure to a third party "substantially increases the opportunity for potential adversaries to obtain the information." *See Merril Lynch & Co., Inc.* v. *Allegheny Energy, Inc.*, 229 F.R.D. 441, 445–46 (S.D.N.Y. 2004).

[40] *See Universal Standard Inc.* v. *Target Corp.*, 331 F.R.D. 80, 88–91 (S.D.N.Y. 2019).

* * *

This memorandum is not intended to provide legal advice, and no legal or business decision should be based on its content. Questions concerning issues addressed in this memorandum should be directed to:

H. Christopher Boehning
+1-212-373-3061
cboehning@paulweiss.com

John P. Carlin
+1-202-223-7372
jcarlin@paulweiss.com

Katherine B. Forrest
+1-212-373-3195
kforrest@paulweiss.com

Elizabeth Hanft
+1-212-373-3664
ehanft@paulweiss.com

David K. Kessler
+1-212-373-3614
dkessler@paulweiss.com

Ian C. Richardson
+1-202-223-7405
irichardson@paulweiss.com

Ross M. Gotler
+1-212-373-2979
rgotler@paulweiss.com

Associates Scott Caravello, Rachel Gallagher, Corey Goldstein, Patrick Lim and Nicholas Tilmes contributed to this Client Memorandum.