

# Update: Judge Rakoff Issues Written Opinion That AI-Generated Documents Are Not Protected by Privilege

February 18, 2026



Last week, [we wrote about a decision](#) in which Judge Rakoff of the Southern District of New York denied the claim of defendant Bradley Heppner that documents prepared by Heppner using the consumer version of the AI model Claude for legal research were privileged. On February 17, 2026, Judge Rakoff issued a written opinion explaining the reasoning behind his February 10 ruling.<sup>1</sup>

## THE COURT'S DECISION

Rakoff accepted that Heppner had (a) been communicating with Claude about factual and legal issues in his case in anticipation of litigation, (b) incorporated information conveyed to him by his counsel during the course of the representation into his communications with Claude, (c) had intended to share the resulting AI-generated documents with counsel, and (d) did, in fact, share those documents with his counsel, but nonetheless rejected both his attorney-client and work product privilege claims.

### Attorney-Client Privilege

Judge Rakoff reasoned that because Claude is not an attorney, communications between Heppner and Claude cannot satisfy the fundamental requirement that privileged

<sup>1</sup> Memorandum, *United States v. Heppner*, No. 25-cr-00503-JSR (S.D.N.Y. Feb. 17, 2026), Dkt. No. 27.

---

communications occur between a client and counsel, noting that “all ‘recognized privileges’ require . . . ‘a trusting human relationship’” such as that with “a licensed professional who owes fiduciary duties and is subject to discipline.”<sup>2</sup> Rakoff held that no such relationship can exist between a user and an AI platform.

Even if the contents of Heppner’s communication with Claude were somehow privileged, Judge Rakoff concluded that the privilege claim would fail because Heppner used the consumer version of Claude, which trains on user data and provides in its privacy policy that Anthropic reserves the right to disclose data to “‘third parties,’ including ‘governmental regulatory authorities.’”<sup>3</sup> Because “[t]he policy clearly puts Claude’s users on notice that Anthropic” may disclose user data “in connection with claims, disputes, or litigation”<sup>4</sup> even absent a subpoena, the Court concluded that Heppner “could have had no ‘reasonable expectation of confidentiality in his communications’ with Claude.”<sup>5</sup> The Court then distinguished the AI-generated exchanges with the consumer version of Claude from confidential notes that a client prepares to share with counsel because “Heppner first shared the equivalent of his notes with a third-party, Claude.”<sup>6</sup>

What is unclear from the decision is whether privilege can never attach to communications made with a consumer AI tool, or whether Heppner merely failed to meet his burden of showing that he had a “reasonable expectation of confidentiality” with Claude.<sup>7</sup> For example, suppose someone like Heppner had uploaded clearly privileged materials into a consumer version of an AI tool and was able to establish that (a) she never read the applicable privacy policy, (b) she did not know that the AI model could train on her data or could share it with third parties, and (c) as a factual matter, considering the millions of queries made to the AI model each day and the nature of the model training, it was extremely unlikely that any human would ever see her communications with the model. The Court’s decision does not foreclose the possibility that the privilege would endure, although the factual burden of establishing (c) may be very difficult to achieve.

What is clearer (as we discussed in our original blog post) is that a court may well take a different view of the reasonable expectation of confidentiality for the use of enterprise AI tools, which do not train on user data and generally provide in their terms and

---

<sup>2</sup> *Id.* at 6 (quoting Ira P. Robbins, *Against an AI Privilege*, JOLT Dig., Harvard L. Sch. (Nov. 7, 2025), <https://jolt.law.harvard.edu/digest/against-an-ai-privilege>).

<sup>3</sup> *Id.* (quoting Anthropic, Privacy Policy (as of February 19, 2025), <https://www.anthropic.com/legal/archive/a2eecf43-807a-4a53-89dd-04c44c351138>).

<sup>4</sup> *Id.* (quoting Anthropic, Privacy Policy).

<sup>5</sup> *Id.* at 7 (quoting *United States v. Mejia*, 655 F.3d 126, 132–34 (2d Cir. 2011)).

<sup>6</sup> *Id.*

<sup>7</sup> *Id.* (quoting *Mejia*, 655 F.3d at 132–34).

---

conditions that user data will not be shared with third parties absent extraordinary circumstances. It is noteworthy that Judge Stein's decision in *In re OpenAI Inc., Copyright Infringement Litigation*, which Judge Rakoff cites for the finding that AI users do not have an expectation of privacy, also involved the use of consumer (as opposed to enterprise) AI tools.<sup>8</sup>

Judge Rakoff ruled that even if Heppner was communicating with Claude for the express purpose of talking to counsel, that was insufficient to establish privilege because Heppner was acting "of his own volition" rather than at counsel's direction,<sup>9</sup> noting that Claude itself expressly disclaims that it can provide legal advice.<sup>10</sup>

The Court noted that had counsel directed Heppner to use Claude, he may have been able to invoke the *Kovel* doctrine, which can extend attorney-client privilege to non-lawyer professionals, such as accountants or consultants, hired by an attorney to assist in providing legal advice. Judge Rakoff stated that, assuming there had been confidentiality and Heppner was acting on instructions from his counsel, in that case "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."<sup>11</sup>

### Attorney Work Product Doctrine

Judge Rakoff dismissed Heppner's work product claims, holding that even assuming the documents were prepared "in anticipation of litigation,"<sup>12</sup> they were not "prepared by or at the behest of counsel"<sup>13</sup> and did not "reflect defense counsel's strategy."<sup>14</sup> The Court declined to follow a magistrate judge's decision in *Shih v. Petal Card, Inc.*,<sup>15</sup> which had extended work product protection to materials prepared by a client without attorney direction, reasoning that such an expansion "undermines the policy animating the work product doctrine," which is to "protect lawyers' mental processes."<sup>16</sup>

---

<sup>8</sup> *In re OpenAI, Inc., Copyright Infringement Litig.*, No. 25-MD-3143 (SHS) (OTW), 2026 WL 21676, at \*1 (S.D.N.Y. Jan. 5, 2026) ("Plaintiffs moved to compel OpenAI to produce a sample of 120 million *consumer* ChatGPT logs") (emphasis added) (cited in Memorandum at 6–7, *Heppner*, No. 25-cr-00503-JSR).

<sup>9</sup> Memorandum at 7, *Heppner*, No. 25-cr-00503-JSR.

<sup>10</sup> *Id.* at 8.

<sup>11</sup> *Id.* at 7 (citing *United States v. Adlman*, 68 F.3d 1495, 1498–99 (2d Cir. 1995); *United States v. Kovel*, 296 F.2d 918 (2d Cir. 1961)).

<sup>12</sup> *Id.* at 9 (quoting *In re Grand Jury Subpoenas Dated March 19, 2002, and August 2, 2002*, 318 F.3d 379, 383 (2d Cir. 2003)).

<sup>13</sup> *Id.* (quoting *In re Grand Jury Subpoenas*, 318 F.3d at 383).

<sup>14</sup> *Id.* (quoting *Matter of Grand Jury Subpoenas Dated Oct. 22, 1991, and Nov. 1, 1991*, 959 F.2d 1158, 1167 (2d Cir. 1992)).

<sup>15</sup> 565 F. Supp. 3d 557 (S.D.N.Y. 2021).

<sup>16</sup> Memorandum at 11, *Heppner*, No. 25-cr-00503-JSR.

---

**KEY TAKEAWAYS**

Judge Rakoff's written opinion reinforces the takeaways we shared last week:

- When using AI tools in connection with privileged communications or legal work, use an enterprise AI tool (which does not train on inputs and maintains the confidentiality of inputs) whenever possible.
- To the extent a client or other non-lawyer is acting at the direction of counsel when using AI tools, accurately document that context in the prompt (e.g., "I am doing this research at the direction of counsel for [X] litigation").
- Privilege logs reflecting such communications should clearly and accurately denote the basis for the privilege and that the AI tool was used with the expectation of confidentiality.

To subscribe to the Data Blog, please [click here](#).

[The Debevoise STAAR \(Suite of Tools for Assessing AI Risk\)](#) is a monthly subscription service that provides Debevoise clients with an online suite of tools to help them fast-track their AI adoption. Please contact us at [STAARinfo@debevoise.com](mailto:STAARinfo@debevoise.com) for more information.

The cover art for this blog post was generated by Chat GPT 5.2.

\* \* \*

Please do not hesitate to contact us with any questions.



**Charu A. Chandrasekhar**  
Partner, New York  
+1 212 909 6774  
[cchandrasekhar@debevoise.com](mailto:cchandrasekhar@debevoise.com)



**Avi Gesser**  
Partner, New York  
+1 212 909 6577  
[agesser@debevoise.com](mailto:agesser@debevoise.com)



**Caroline H. Wallace**  
Associate, New York  
+1 212 909 6777  
[chwallace@debevoise.com](mailto:chwallace@debevoise.com)

*This publication is for general information purposes only. It is not intended to provide, nor is it to be used as, a substitute for legal advice. In some jurisdictions it may be considered attorney advertising.*