
SDNY Judge Finds Criminal Defendant's Case-Strategy Communications with Gen-AI Tool Not Privileged or Work Product

In *United States v. Heppner*, No. 1:25-cr-00503-JSR, 2026 U.S. Dist. LEXIS 32697 (S.D.N.Y. Feb. 17, 2026), United States District Court Judge Jed S. Rakoff granted the government's motion for a ruling that a criminal defendant's (Bradley Heppner) written exchanges about case strategy with a public generative artificial-intelligence ("gen-AI") platform (Claude) were not protected by the attorney-client privilege or the attorney work-product doctrine. Specifically, the court granted the government's motion for access to case-strategy documents created by the criminal defendant using Claude before his arrest on federal fraud charges.

This decision, addressing a question of first impression, comes at a time when the use of gen-AI tools has become ubiquitous in business and social settings. It sets an important precedent to be further considered or distinguished by the courts, with these key takeaways:

- Heppner's communications with the gen-AI tool were not communications with an attorney, therefore they could not be covered by the attorney-client privilege.
- Heppner's communications with the gen-AI tool were not confidential, because Claude's privacy policy states that it collects users' inputs and outputs for training purposes and reserves the right to disclose them voluntarily to third parties.
- Heppner did not communicate with the gen-AI tool primarily to obtain legal advice, because he did not interact with the gen-AI tool at the suggestion or direction of counsel.
- Heppner prepared the AI-generated documents of his own volition, not at the behest of counsel, therefore the documents were not protected as attorney work product under Second Circuit precedent.

This decision highlights the importance of carefully regulating the use of gen-AI tools in the legal context.

Background

Bradley Heppner was indicted by a grand jury in New York on October 28, 2025, for securities fraud, wire fraud, conspiracy to commit securities fraud and wire fraud, making false statements to auditors, and falsifying corporate records, in connection with his role as executive of several corporate entities. Heppner was arrested by the FBI on November 4, 2025. Pursuant to a search warrant, the FBI seized from his home 31 documents memorializing communications that Heppner had with Claude, a public gen-AI tool offered by Anthropic. The documents set out Heppner's research into defense strategy after his earlier grand jury subpoena and discussions with the government regarding the criminal investigation.

Heppner's legal counsel asserted privilege over the documents and averred that the documents were created for the purpose of later speaking with counsel to obtain legal advice and were subsequently shared with counsel, although they were not created at the direction of counsel. The government filed a motion on February 6, 2026, seeking access to the documents and a ruling that the documents were not protected by privilege or work-product doctrine.

After oral argument on February 10, 2026, the court granted the motion. Its written decision issued on February 17, 2026.

I. The SDNY Decision

a. Attorney-Client Privilege

The court cited the longstanding principle that attorney-client privilege protects from disclosure communications: (1) between client and attorney; (2) kept confidential; and (3) made for the purpose of obtaining or providing legal advice.¹ The court found that “the AI Documents lack at least two, if not all three, elements of the attorney-client privilege.”²

The court noted that Heppner did not and could not allege that the public gen-AI tool, Claude, was an attorney, thus communications with Claude were not communications with an attorney. The court concluded that “[b]ecause Claude is not an attorney [...] that alone disposes of Heppner’s claim of privilege.”³ Nevertheless, the court further discussed its awareness of commentators’ views that gen-AI tools such as Claude should not be treated differently from “the use of other Internet-based software, such as cloud-based word processing applications.”⁴ The court then quoted a *Harvard Journal of Law and Technology* article, titled “Against an AI Privilege,”⁵ for the proposition that all privileges require “a trusting human relationship,” in this case with an attorney. The court therefore dismissed the possibility of such a relationship existing between a human user and the public gen-AI platform.

The court’s discussion does not clearly address whether a pre-existing human relationship with an attorney would yield a different result, for instance if a client’s use of a gen-AI tool was directed or suggested by its legal counsel. Significantly, about the same time as the *Heppner* decision, another federal district court declined, in a civil matter (*Warner v. Gilbarco, Inc.*, No. 2:24-cv-12333, 2026 U.S. Dist. LEXIS 27355 (E.D. Mich. Feb. 10, 2026)), to compel discovery from a *pro se* plaintiff of “all documents and information concerning [plaintiff’s] use of third-party AI tools in connection with this lawsuit.”⁶ That decision focused on the work-product doctrine and made clear that, because the plaintiff was *pro se*, plaintiff’s interactions with the AI tool were protected by the work product doctrine, in part because plaintiff was representing herself and in part because her use of AI did not make it more likely that plaintiff’s adversary would get access to the materials at issue.⁷

On the second prong, confidentiality of the communications, the *Heppner* court found that the use of Claude, provided by the Anthropic company, lacked confidentiality due to the defendant’s consent to Anthropic’s privacy policy. The privacy policy sets out that Claude collects data from users’ inputs and Claude’s outputs and uses the data to train the gen-AI tool. The court did not address the possibility, offered by Claude and other gen-AI tools, of opting out of model training (a notable option for all paid tiers of Claude accounts) and the enhanced confidentiality that such opt-outs may provide. The court further considered, however, that Anthropic’s privacy policy enables it to

¹ *Heppner*, 2026 U.S. Dist. LEXIS 32697, at *4 (citing *United States v. Mejia*, 655 F.3d 126, 132 (2d Cir. 2011)).

² *Id.*, at *5.

³ *Id.*

⁴ *Id.*

⁵ *Id.*, at *6 (citing 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>).

⁶ *Warner v. Gilbarco, Inc.*, 2026 U.S. Dist. LEXIS 27355, at *11.

⁷ *Id.* at *11-12.

share customer data voluntarily with “a ‘host of third parties,’ including ‘governmental regulatory agencies,’” including “in connection with claims, disputes or litigation.”⁸ The court found that Heppner’s documents had been shared with a third party with no reasonable expectation of confidentiality.

On this point, it is worth noting that the confidentiality policies of gen-AI providers vary considerably among providers and among a given provider’s various offerings. Several providers provide tiers of gen-AI tools, with enterprise or professional tiers that may be distinguishable by their prohibition of collection and sharing of user inputs, outputs, and other customer data. Use of tools with tighter confidentiality provisions will make it more likely that privilege assertions will be upheld.

For the third prong, the court admitted that the issue whether Heppner communicated with Claude for the purpose of obtaining legal advice “perhaps presents a closer call”⁹ because Heppner’s counsel asserted that this was the purpose of those communications. The court concluded, however, that this condition was not satisfied because Heppner did not communicate with Claude “at the suggestion or direction of counsel.”¹⁰ Heppner’s later sharing of those communications with counsel did not transform them into privileged communications.¹¹ The court noted that the gen-AI tool could have acted similarly to the attorney’s agent, had Heppner used it at the direction of counsel. The court rejected the argument that Heppner’s sharing of the materials with counsel later evidenced an intent to generate the materials for the purpose of obtaining legal advice. The court’s holding highlights the importance of setting up, beforehand, a framework for counsel to direct the use of gen-AI tools when such use may be relevant to litigation, so as to maximize the chances that privilege and/or work product protections will apply.

b. Attorney Work-Product Doctrine

The court also held that Heppner’s communications with the public gen-AI tool did not constitute attorney work product. Citing Second Circuit precedent,¹² the court considered that while the communications may have been prepared in anticipation of litigation, the work-product doctrine does not attach to them because they were not “prepared by or at the behest of counsel [...] nor did they reflect defense counsel’s strategy.”¹³ The court dismissed Heppner’s argument that Federal Rule of Criminal Procedure 16(b)(2)(A) excluded pre-trial discovery of “reports, memoranda or other documents made by the defendant, or the defendant’s attorney or agent, during the case’s investigation or defense” because the communications with the public gen-AI tool were seized pursuant to a valid search warrant (instead of produced in response to a discovery request).¹⁴ The court also rejected Heppner’s argument that the court should follow the decision in *Shih v. Petal Card, Inc.*, 565 F. Supp. 3d 557 (S.D.N.Y. 2021), which held that the work product doctrine was not limited to materials prepared by or at the direction of an attorney (provided they were prepared by a defendant in preparation of litigation and to seek legal counsel). The court

⁸ *Heppner*, 2026 U.S. Dist. LEXIS 32697, at *6.

⁹ *Id.*, at *7.

¹⁰ *Id.*

¹¹ *Id.*, at *8.

¹² *In re Grand Jury Subpoenas Dated March 19, 2002, and August 2, 2002*, 318 F.3d 379, 383 (2d Cir.2003), *Matter of Grand Jury Subpoenas Dated Oct. 22, 1991, and Nov. 1, 1991*, 959 F.2d 1158 (2d Cir. 1992).

¹³ *Heppner*, 2026 U.S. Dist. LEXIS 32697, at *9.

¹⁴ *Id.*, at *10.

disagreed with this holding as “undermin[ing] the policy animating the work product doctrine,”¹⁵ whose purpose “is to protect lawyer’s mental processes.”¹⁶

Heppner’s narrow interpretation of the work-product doctrine may lead courts to distinguish it when following precedent or applicable criminal or civil rules that protect documents prepared by parties in anticipation of litigation under the attorney work-product doctrine. For instance, Fed. R. Civ. P. 26(b)(3)(A) states that “[o]rdinarily, a party may not discover documents and tangible things that are prepared in anticipation of litigation or for trial by another party or its representative.” As noted above, in a written decision issued on the same day as the oral argument in *Heppner*, the E.D. Michigan court in *Warner v. Gilbarco* found that a civil *pro-se* party’s gen-AI documents prepared in anticipation of litigation were shielded under Fed. R. Civ. P. 26(b)(3)(A). The *Warner* court held that the use of a gen-AI tool did not waive the work-product doctrine as “to the extent Defendants argue that Plaintiff waived the work-product protection by using ChatGPT, the work-product waiver has to be a waiver to an adversary or in a way likely to get in an adversary’s hand.”¹⁷

II. Implications

The *Heppner* decision is significant because courts have only recently begun to assess the impact of the use of public gen-AI on the attorney-client privilege and the attorney-work product doctrine. While we expect that further cases will provide clearer guidance about the application of the attorney-client privilege and work-product doctrine to gen-AI communications, *Heppner*’s broad holding denying such protections signals the importance of carefully crafting gen-AI policies and advising clients about the creation and use of such materials. Such policies should ensure the use of gen-AI, in matters of relevance to prospective or current litigation, is limited to gen-AI tools with robust confidentiality provisions (e.g., that do not share inputs or outputs with third parties or use such data for training purposes) and is directed by legal counsel.

* * *

If you have any questions about the issues addressed in this alert, or if you would like a copy of any of the materials mentioned in it, please do not hesitate to call or email authors Joel Kurtzberg (partner) at 212.701.3120 or jkurtzberg@cahill.com; Britton Davis (partner) at 202.862.8932 or bdavis@cahill.com; or Raphael Chabaneix (associate) at 202.862.8972 or rchabaneix@cahill.com; or email publicationscommittee@cahill.com.

¹⁵ *Id.*, at *11.

¹⁶ *Id.*

¹⁷ *Warner*, 2026 U.S. Dist. LEXIS 27355, at *12.