ALERT MEMORANDUM

Courts Grapple with Privilege Implications of AI

February 27, 2026

In February 2026, two federal courts issued first-of-their-kind decisions on a previously untested question: whether materials generated using publicly available AI tools are protected from discovery. Judge Rakoff of the United States District Court for the Southern District of New York held that documents a criminal defendant generated using a publicly available AI tool were not privileged or protected work product. Magistrate Judge Anthony Patti of the Eastern District of Michigan reached a different conclusion, holding that a *pro se* civil litigant's ChatGPT queries and the AI's responses were protected by the work-product doctrine. These are significant decisions that will shape the developing law in this area as courts, litigants, and practitioners navigate the intersection of artificial intelligence and established privilege doctrines.

If you have any questions concerning this memorandum, please reach out to your regular firm contact or the following authors

NEW YORK

Rahul Mukhi
+1 212 225 2921
rmukhi@cgsh.com

Jon Kolodner
+1 212 225 2690
jkolodner@cgsh.com

Arminda Bepko
+1 212 225 2517
abepko@cgsh.com

Nathaniel (Nate) Reynolds
+1 212 225 2889
nreynolds@cgsh.com

SILICON VALLEY

Angela Dunning
+1 650 815 4131
adunning@cgsh.com

WASHINGTON

David A. Last
+1 202 974 1650
dlast@cgsh.com

United States v. Heppner

On February 17, 2026, Judge Rakoff of the United States District Court for the Southern District of New York held that documents a criminal defendant generated using a publicly available AI tool were not privileged or protected work product.[1] The Court found that, under the circumstances presented—including that the defendant did not act at counsel's direction and was informed that his data might be disclosed to third parties—neither doctrine applied, and held that the documents were not protected.

On November 5, 2025, the FBI arrested defendant Heppner and seized his electronic devices.[2] While under government investigation—after receiving a grand jury subpoena—Heppner had previously run several queries through the publicly available version of Claude.[3] Heppner's queries and Claude's responses (collectively, the "AI Documents") were located on his electronic devices when the government executed its search warrant. Heppner later shared the AI Documents with defense counsel, who conceded that Heppner created the AI Documents of his own volition, without direction from counsel.

In the criminal prosecution, the defendant sought to preclude the government from reviewing the AI Documents on the grounds of attorney-client privilege and work-product protection. Following a bench ruling on February 10, 2026, Judge Rakoff issued a written opinion on February 17.

Attorney-Client Privilege

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] Judge Rakoff determined that the AI Documents lacked "at least two, if not all three, elements of the attorney-client privilege."[5]

First, Judge Rakoff observed that the AI Documents are not communications between client and attorney.[6] Rather, the communications were between Heppner and Claude, which "is not an attorney."[7] Judge Rakoff emphasized that for any recognized privilege to apply, there must be a "trusting human relationship." Because "no such relationship exists" between a user and an AI tool, "that alone disposes of Heppner's claim of privilege."[8]

Second, Heppner failed to establish that his communications with Claude were intended to be, and were kept, confidential. The Court cited *In re OpenAI, Inc., Copyright Infringement Litig*ation, in which the Southern District of New York observed that AI users do not have "substantial privacy interests" in their communications with "publicly accessible" AI platforms.[9] Judge Rakoff also noted that Claude's privacy policy explicitly states that Anthropic reserves the right to disclose user data to third parties, including "governmental regulatory authorities."

Third, Heppner did not satisfy the requirement that the communication be made for the purpose of obtaining or providing legal advice. The defendant conceded that he did not use Claude at the direction of counsel. Judge Rakoff suggested that if defense counsel had directed Heppner to use AI, "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."[10] But that is not what occurred here. The Court also observed that when asked whether it can provide legal advice, Claude responded: "I'm not a lawyer and can't provide formal

[1] *U.S. v. Heppner*, No. 25-CR-503, ECF No. 27 at 2 (S.D.N.Y. Feb. 17, 2026).
[2] *Id.*
[3] *Id.* at 3.
[4] *Id.* at 4 (citing *U.S. v. Mejia*, 655 F.3d 126, 132 (2d Cir. 2011)).
[5] *Id.* at 5.
[6] *Id.* (citing *In re OpenAI, Inc., Copyright Infringement Litig.*, 802 F. Supp. 3d 688, 699 (S.D.N.Y. 2025)).
[7] *Id.* at 5.
[8] *Heppner*, No. 25-CR-503, ECF No. 27 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.).
[9] *Id.* at 6.
[10] *Id.* at 7. (citing *United States v. Adlman*, 68 F.3d 1495, 1498-99 (2d Cir. 1995)).

legal advice or recommendations," and recommended that the user consult with a "qualified attorney."[11]

Regardless of whether Heppner intended to share the AI Documents with counsel, Judge Rakoff emphasized that it is "black-letter law that non-privileged communications are not somehow alchemically changed into privileged ones upon being shared with counsel."[12]

In a footnote, Judge Rakoff observed that even if the AI Documents were privileged because they incorporated information from counsel, Heppner waived any such privilege by using Claude. By sharing the information with Claude and Anthropic, Heppner waived privilege "just as if he had shared it with any other third party."[13] However, the factual record on the waiver issue was limited. For example, the Court did not address whether Heppner was preparing a communication for his attorney, what the terms of service provide regarding confidentiality, how prominent those disclosures were, whether Heppner opted out of allowing his material to be used for training (as Claude permits), or, if he did not opt out, whether any human at Anthropic would ever review the material.

Work-Product Doctrine

The Court also held that the AI Documents were not protected by the work-product doctrine, which shields from disclosure "materials prepared by or at the behest of counsel in anticipation of litigation or for trial."[14] Judge Rakoff emphasized that "at its core," the doctrine protects *attorneys*' mental processes and provides "a privileged area within which [counsel] can analyze and prepare his client's case."[15] The Second Circuit has repeatedly held that the doctrine's purpose "is not generally promoted by shielding from discovery materials in an attorney's possession that were prepared neither by the attorney nor his agents.[16] Applying these principles, the Court identified two fatal deficiencies: first, defense counsel confirmed that the AI Documents "were prepared by the defendant on his own volition," meaning Heppner was not acting at the direction of his lawyers when he communicated with Claude; and second, while counsel plausibly asserted that the AI Documents "affected" the defense's strategy going forward, counsel conceded that the documents did not "reflect" counsel's strategy at the time Heppner created them.[17]

The Court rejected Heppner's reliance on Federal Rule of Criminal Procedure 16(b)(2)(A), which provides that when a criminal defendant responds to pretrial discovery requests, the defendant need not produce "documents made by the defendant, or the defendant's attorney or agent, during the case's investigation or defense."[18] Judge Rakoff found this rule inapplicable because the AI Documents were seized pursuant to a search warrant executed at the time of Heppner's arrest—the government did not request them, and Heppner did not produce them in pretrial discovery. Because the rule applies only to materials a defendant would otherwise be obligated to produce in response to government discovery requests, it was inapplicable on its face.

The Court also declined to follow *Shih v. Petal Card, Inc.*, in which a Magistrate Judge held that the work-product doctrine protected certain communications a plaintiff prepared in anticipation of litigation "regardless of whether [her lawyer] was acting as her counsel at the time, and without showing that another attorney 'directed the work.'"[19] Judge Rakoff disagreed with *Shih*, noting that it was not binding and reasoning that extending work-product protection to materials not prepared at counsel's direction "undermines the policy animating the work-product doctrine," which is "to

[11] *Heppner,* ECF No. 23-6 at 1-2.
[12] *Heppner*, ECF No. 27 at 8.
[13] *See United States v. Jacobs*, 117 F.3d 82, 91 (2d Cir. 1997) (stating that voluntary disclosure of confidential material to a third party waives any applicable attorney-client privilege).
[14] *In re Grand Jury Subpoenas Dated March 19, 2002 and August 2, 2022*, 318 F.3d 379, 383 (2d Cir. 2003).
[15] *Heppner*, ECF No. 27 at 8-9 (quoting *U.S. v. Noble*s, 422 U.S. 225, 238 (1975)) (emphasis added).
[16] *Id.* at 8,10–11.
[17] *Id.* at 9–10.
[18] *Id.* at 10.
[19] *Id.* at 10–11.

preserve a zone of privacy in which a lawyer can prepare and develop legal theories and strategy 'with an eye toward litigation.'"[20] While the work-product doctrine may apply to materials generated by non-lawyers, the Second Circuit has consistently emphasized that the doctrine's purpose is to protect lawyers' mental processes—not to shield any document a client creates with an eye toward eventual litigation.

Judge Rakoff concluded that because the AI Documents were not prepared at the behest of counsel and did not disclose counsel's strategy, they do not merit protection as work product.

Warner v. Gilbarco, Inc.

On February 10, 2026, Magistrate Judge Anthony Patti of the Eastern District of Michigan held that a *pro se* civil litigant's ChatGPT queries and the AI's responses are protected by the work-product doctrine.[21] In this ruling, Magistrate Judge Patti found that the AI-related materials reflected the plaintiff's mental impressions prepared in anticipation of litigation. The Court further held that, under Sixth Circuit law, disclosing information to ChatGPT did not constitute a waiver of work-product protection because such waiver requires disclosure to an adversary or in a manner likely to reach an adversary's hands.

Plaintiff Warner, proceeding *pro se*, brought suit against her former employer Gilbarco, Inc., in a civil employment dispute.[22] After Warner admitted to using ChatGPT to answer legal questions and draft filings, the defendants moved to compel production of plaintiff's queries and ChatGPT's responses (the "AI Materials").[23] The plaintiff asserted work-product protection in response.[24]

The Court in *Warner* found that the work-product doctrine protected the AI Materials because: (1) they reflected the thought processes of a *pro se* litigant and (2) the plaintiff did not waive work-product protection by divulging information to ChatGPT.[25] Under Sixth Circuit law, which is consistent with law in other jurisdictions, work product waiver occurs only when disclosure is made to an adversary or in a way likely to place the materials in an adversary's hands.[26] (This is distinguished from attorney-client privilege waiver which can occur even if disclosure is not to an adversary). In *Warner*, the court held that the plaintiff's disclosure of information to ChatGPT was not disclosure to an adversary; it was disclosure to an AI software. And as Judge Patti noted: AI programs "are tools, not persons, even if they may have administrators somewhere in the background."[27] Agreeing with plaintiff, the Court observed that compelling production here would force disclosure of a litigant's internal mental impressions, and accepting defendants' theory "would nullify work-product protection in nearly every modern drafting environment[.]"[28]

While the *Heppner* and *Warner* courts arrived at opposite conclusions about whether information shared with generative AI tools should be shielded from discovery, each case turns on its own facts. In *Heppner*, defense counsel conceded that the defendant created the AI Documents of his own volition—without counsel's direction—and that the AI Documents did not reflect counsel's mental impressions or litigation strategy at the time they were created. In *Warner*, however, the *pro se* plaintiff's AI Materials reflected her own mental impressions and litigation strategy developed in anticipation of litigation, and the Court held they were protected work product. While both cases arise in novel legal and technological contexts, neither case disturbs

[20] *Id.* at 11.(citations omitted).
[21] *Warner v. Gilbarco, Inc.*, No. 2:24-cv-12333, ECF No. 94 at 10-12 (E.D. Mich. Feb. 10, 2026).
[22] *Id.* at 1–2.
[23] Defs.' Mot. to Compel and Request for Sanctions, *Warner v. Gilbarco*, No. 2:24-cv-12333, ECF No. 80 at 2–3 (E.D. Mich. Dec. 23, 2025).
[24] Pl.'s Opp. to Defs.' Mot. to Compel and Request for Sanctions, *Warner v. Gilbarco*, No. 2:24-cv-12333, ECF No. 86 at 4 (E.D. Mich. Jan. 6, 2026).
[25] *Warner v. Gilbarco, Inc.*, No. 2:24-cv-12333, ECF No. 94 at 10–12 (E.D. Mich. Feb. 10, 2026).
[26] *Id.* at 11(citing In re Columbia/HCA Healthcare Corp. Billing Pracs. Litig., 293 F.3d 289, 306 n.28 (6th Cir. 2002).
[27] *Warner*, ECF No. 94 at 12.
[28] *Id.* at 12–13.

the fundamental understanding that the work-product doctrine exists to protect attorneys' mental processes.

Key Takeaways

— **The law is unsettled and still developing.** *Heppner* and *Warner* are the first two decisions to address whether AI-generated materials are protected from discovery, and they reached opposite conclusions. Both decisions may face appellate scrutiny, and additional courts will inevitably confront these issues. In the interim, companies and in-house counsel should carefully evaluate their own AI usage practices in light of these decisions.

— **Both rulings were fact-driven and leave questions unresolved**. Each court's analysis was heavily influenced by the specific circumstances before it—including the nature of the litigation, the party's representation status, and the AI tool's privacy terms. Neither decision addresses whether a different outcome might result if a represented party uses AI at counsel's direction, or if a party uses an enterprise AI tool with negotiated confidentiality protections. A closed AI system that does not permit disclosure to third parties should be readily distinguishable from the publicly available tools at issue in both cases. Nor does either decision address tools that are used in a "safe" or "confidential" mode in which inputs and outputs are maintained in confidence and not used to train the models.

— **Courts are likely to focus on whether AI-generated materials were created at counsel's direction and reflect litigation strategy**. In *Heppner*, defense counsel conceded that the AI Documents were created without counsel's direction and did not reflect counsel's strategy—and protection was denied. In *Warner*, as a *pro se* plaintiff, Warner's AI Materials reflected her own litigation strategy, and protection was granted. Companies should ensure that any use of AI in connection with legal advice is directed by counsel and documented accordingly.

— **Sharing information with publicly available AI tools carries risk**. While *Warner* held that using ChatGPT did not waive work-product protection, *Heppner* emphasized that Anthropic's privacy policy expressly disclaims any expectation of privacy in user inputs—and treated disclosure to a public AI tool as equivalent to disclosure to a third party for privilege purposes. Before inputting any privileged, confidential, or litigation-related information into an AI tool, users should consult with counsel and carefully review the tool's privacy policies and terms of service. The risk of disclosure is highest with publicly available tools that disclose potential sharing of information with third parties, including where the information could be used to train the tool being used.

— **AI tools are not attorneys, and communications with them are not privileged absent counsel's direction**. Judge Rakoff's decision in *Heppner* makes clear that an AI tool—no matter how sophisticated—is not an attorney. Communications with AI platforms do not satisfy the elements of attorney-client privilege when they are not made at counsel's direction. Litigants should not assume that sharing AI-generated materials with counsel creates privilege, particularly if the use of AI was not directed by counsel.

— **Firms and clients should adopt AI-usage policies**. Clients should use AI in legal matters only at counsel's direction, and attorneys should review AI platform privacy policies before inputting sensitive information. Companies that use AI tools in legal work should adopt internal AI-usage policies governing the handling of privileged and confidential information, and document how those tools are being used.

…

This article was prepared with contributions from Cleary law clerks, Rebecca Quillen and Emma Valiante.

CLEARY GOTTLIEB