The Chatbot Confessional

By **Jonathan Franklin**, Shareholder | March 9, 2026



A Man in Trouble

Bradley Heppner was under indictment in the Southern District of New York when he sat down with Anthropic's Claude. Over several sessions, he prepared roughly thirty-one documents — reports outlining defense strategy, arguments about facts and law, analysis of what the government might charge. He did this on his own. His lawyer did not ask him to do this.

That fact turned out to matter more than anything Heppner typed into the chat window. On February 17, 2026, Judge Jed Rakoff ruled that the documents were protected by neither attorney-client privilege nor the work product doctrine. The government could read every word. *United States v. Heppner*, No. 1:25-cr-00503-JSR (S.D.N.Y. Feb. 17, 2026).

One week earlier and six hundred miles west, a *pro se* plaintiff in the Eastern District of Michigan had used ChatGPT to prepare her litigation materials. The defendants in *Warner v. Gilbarco, Inc.*, No. 2:24-cv-12333 (E.D. Mich. Feb. 10, 2026), moved to compel everything related to her AI use. Magistrate Judge Anthony Patti denied the motion, concluded the materials were protected work product, and called the pursuit of them "a distraction from the merits of this case." A pro se plaintiff, representing herself with no attorney involved, kept her AI work product. Heppner, who had counsel, lost his AI documents entirely.

The commentary has treated these opinions as the opening chapter of a new AI privilege doctrine. They are narrower than that, more fact-bound, and more platform-specific. This piece is about the fine print.

The Fine Print

Start with what Rakoff actually cited. The *Heppner* opinion rests a significant part of its confidentiality analysis on Anthropic's privacy policy. The critical language: Anthropic "reserves the right to disclose such data to a host of 'third parties,' including 'governmental regulatory authorities.'" That provision appears in both the 2025 and 2026 versions of the policy, and it creates a serious problem for any confidentiality claim. A user who accepts terms permitting disclosure to governmental authorities faces a steep climb in arguing the communication remained confidential.

Rakoff relied on several grounds together — no attorney, no confidentiality, no legal purpose — but the disclosure language may do much of the work on its own. Either way, the terms governing the engagement are not uniform, and this is where the fine print gets interesting.

Anthropic's commercial terms of service — governing Team, Enterprise, and API accounts — prohibit training on customer content, impose contractual confidentiality obligations, and require notification before compelled disclosure. The consumer terms that likely governed Heppner's sessions contain none of those protections. Same platform. Same model. A different contractual architecture on every element Rakoff relied on.

A separate ruling helps mark the baseline. In the OpenAI copyright MDL, Judge Stein affirmed an order compelling production of twenty million de-identified consumer ChatGPT logs, holding that "privacy interests in users' conversations with ChatGPT which users voluntarily disclosed to OpenAI and which OpenAI retains in the normal course of its business" are weaker than the privacy interests in wiretapped phone conversations. *In re OpenAI, Inc. Copyright Infringement Litig.*, No. 25-md-3143, Dkt. 1021 (S.D.N.Y. Jan. 5, 2026). That was a discovery ruling, not a privilege holding — but Rakoff himself cited it in *Heppner* for the proposition that AI users lack substantial privacy interests in platform conversations. Discoverable consumer chat logs are not the same thing as privileged legal communications or protected work product. The OpenAI order shows what the unprotected end of the spectrum looks like.

Old Rules, New Machines

If these district court opinions are data points, where are the benchmarks? The Supreme Court has not decided whether AI interactions are privileged. But its existing cases already identify every variable the lower courts are working through.

In *Upjohn Co. v. United States*, 449 U.S. 383 (1981), the Court tied privilege to confidential communications made for the purpose of obtaining legal advice. It warned that "[a]n uncertain privilege, or one which purports to be certain but results in widely varying applications by the courts, is little better than no privilege at all." 449 U.S. at 393. In *United States v. Jicarilla Apache Nation*, 564 U.S. 162 (2011), the Court returned to the point: "for the attorney-client privilege to be effective, it must be predictable." 564 U.S. at 183. A privilege analysis that turns too heavily on opaque platform terms is the opposite of predictable. The existing framework — structured engagement, commercial terms, attorney direction, anticipation of litigation — provides that predictability.

In *United States v. Nobles*, 422 U.S. 225 (1975), the Court recognized that the work product doctrine "shelters the mental processes of the attorney" and extends to "material prepared by agents for the attorney as well as those prepared by the attorney himself." 422 U.S. at 238–39. That is the Supreme Court hook for properly directed AI use — and it is what Rakoff's *Kovel* dictum pointed toward without holding.

Work product deserves particular attention because the doctrinal path may be broader than *Heppner* suggests. Federal Rule of Civil Procedure 26(b)(3)(A) protects materials "prepared in anticipation of litigation or for trial by or for another party or its representative (including the other party's attorney, consultant, surety, indemnitor, insurer, or agent)." Federal Rule of Criminal Procedure 16(b)(2)(A) protects "reports, memoranda, or other documents made by the defendant, or the defendant's attorney or agent, during the case's investigation or defense." Both rules extend beyond attorney-authored materials. Read together, they suggest that attorney direction is not the exclusive gateway to work product protection.

In *Fisher v. United States*, 425 U.S. 391 (1976), the Court drew the limit: privilege protects communications, not underlying facts or preexisting materials. Some AI inputs and outputs will remain unprotected even when incorporated into legal work. The doctrine has limits. Those limits apply to AI the same way they apply to everything else.

The Trouble with Heppner

Heppner is a pyrrhic opinion. It does not help the side that won, and it does not help the side that lost.

Rakoff held that Rule 16(b)(2)(A) was inapplicable on its face because Heppner's documents were seized by warrant, not produced in reciprocal discovery. That procedural holding is narrow. But the broader textual question remains unresolved: Rule 16(b)(2)(A) protects documents made by "the defendant, or the defendant's attorney or agent, during the case's investigation or defense." Rule 26(b)(3) lists attorney as one of several qualifying representatives. The better reading of both rules may be that attorney direction is one way to establish work product protection, not the exclusive gateway — and that the central textual inquiry is anticipation of litigation. Judge Patti in *Warner* suggested as much: he recognized work product protection for a pro se litigant's AI-assisted litigation materials without requiring attorney direction. The differing outcomes may reflect differences in textual emphasis, but also in posture, circuit law, waiver doctrine, and factual setup — criminal case versus civil case, represented defendant versus pro se plaintiff, warrant seizure versus discovery request.

The opinion is also a warning. Consumer AI use without attorney direction is exposed. The risk, under Rakoff's reasoning, is a kind of waiver by input: privileged material may lose its protection when typed into a consumer chatbot outside a protected arrangement. Rakoff's *no alchemy rule* — that "non-privileged communications are not somehow alchemically changed into privileged ones upon being shared with counsel" — means you cannot launder unprotected AI work through counsel after the fact.

These are district court opinions. They bind no other court — not even other judges in the same district. No circuit court has yet spoken. No state supreme court has yet spoken. They are data points, not benchmarks.

What Was Left Open

The question is not whether AI communications can be privileged. The Supreme Court's existing framework appears capable of accommodating them under the right conditions. The question is whether you structured the engagement to satisfy those conditions.

Structure does the work. The strongest case for protection will usually involve counsel selecting the platform, directing the task, and incorporating the system into a matter-specific legal workflow. That also means terms that actually protect confidentiality — contractual restrictions, training prohibitions, and notice before compelled disclosure. Not a consumer-facing service whose terms reserve broad disclosure rights to governmental authorities. And of the two doctrines, work product may prove the more durable shield. It protects counsel's mental processes, extends to agents under *Nobles,* and the textual reach of both Rule 26(b)(3) and Rule 16(b)(2)(A) suggests a broader path than *Heppner* acknowledged.

No court needs to invent a new AI privilege. Lawyers need to structure AI use to satisfy the old requirements they already know.

JONATHAN FRANKLIN
Shareholder, Franklin Legal Group, PA
Jonathan@FranklinLG.com
305.677.8010

Disclaimer: This discussion is commentary, not legal counsel. The court opinions discussed here are public. Read them. Read the rules. And before you type anything sensitive into a chatbot, talk to your lawyer — not the chatbot. The image was generated using Nano Banana 2.