NEW YORK COURT RULES THAT YOUR AI CHATS CAN BECOME EVIDENCE AGAINST YOU

Update by: Pooja Tidke, Rhea Rao | April 06, 2026

On 10th February 2026, in United States v. Heppner ("*Heppner*"), Hon. Jed S. Rakoff of the United States District Court for the Southern District of New York ruled that written exchanges that were had by a party with the consumer version of the generative AI platform 'Claude' were not protected by attorney-client privilege or the 'work product' doctrine (which protects materials prepared by or at the behest of a lawyer). On 17th February 2026, Judge Rakoff issued a written opinion supporting his previous decision. This opinion is demonstrably monumental considering that it highlights the considerations that would be relevant to parties who are not advocates, when using generative AI tools moving forward.

Brief Facts

On 28th October 2025, a federal grand jury returned an indictment charging Bradley Heppner with securities fraud, wire fraud, conspiracy to commit securities and wire fraud, making false statements to auditors and falsification of records. The indictment alleged that Heppner engineered a scheme to defraud a listed entity's investors by making false representations about and causing the publicly traded GWG Holdings Inc. to enter into transactions, with two entitles that Heppner had created and controlled, for his personal benefit.

Heppner was arrested on 4th November 2025 and at the time of arrest the FBI executed a search warrant and seized numerous hard copy records and dozens of electronic devices. Shortly thereafter, his lawyers informed the Government that prior to his arrest and after the receipt of a grand jury subpoena, Heppner used the generative AI platform Claude (operated by Anthropic) to run queries related to the Government's investigation and later shared this material with his lawyers. The resulting documents, approximately thirty-one in number, were expected to be found on the seized devices. These documents comprised reports that outlined defence strategy outlining possible arguments that he could make on facts and the law that he might be charged in relation thereto. Heppner's lawyers asserted privilege over these "AI Documents".

Arguments made by the Parties

Plaintiff (Government)

- The AI Documents fail all three elements of attorney-client privilege: (a) they are not communications between a client and attorney as Claude is not a lawyer; (b) they were not made for the purpose of obtaining legal advice and Claude's own terms and policies explicitly disclaim providing legal advice; and (c) they are not confidential as Heppner voluntarily shared his prompts with a third-party commercial platform whose privacy policy expressly permits disclosure to governmental authorities.
- The subsequent transmission of the AI Documents to his lawyers does not retroactively cloak them with privilege.
- Since the Defendant had conceded that he was not directed by his counsel to run the Claude searches, the work product doctrine was inapplicable. The policy rationale of the doctrine which is to preserve a zone of privacy for the attorney's mental processes, simply does not apply where the defendant independently generated the materials without counsel involvement. Further, Claude has no law degree, is not a member of any bar, owes no fiduciary or professional duties to users and is not subject to professional discipline. Thus, attorney-client privilege cannot be mapped onto a machine.

Defendant (Heppner)

- The AI Documents incorporated information that Heppner's lawyers had themselves conveyed to him and were created for the "*express purpose of talking to counsel*" and obtaining legal advice, after which they were in fact shared with the Defendant's lawyers.
- It was argued that the work-product doctrine does not require attorney direction and that work product protection applies regardless of whether an attorney directed the client's

preparation of the materials. The Defendant's lawyers conceded, however, that they "*did not direct Heppner to run Claude searches*" and while the AI Documents "*affected*" their strategy, they did not "*reflect*" their strategy at the time of preparation.

Court's Decision

I. Attorney-Client Privilege

A. Not Client and Attorney Communication

The Court held that Claude is not an attorney and the AI Documents are therefore not communications between a client and counsel. The Court acknowledged the academic argument that Heppner's AI inputs might be analogised to the use of cloud-based word processing applications rather than communications with a third party but observed that such an argument would in any event cut against his interests. All recognised attorney-client privileges "*require, among other things "a trusting human relationship*," such as, in the attorney-client context, as relationship "*with a licensed professional who owes fiduciary duties and is subject to discipline*." No such relationship can exist between an AI user and a platform such as Claude.

B. Lack of Confidentiality

The Court found the communications were not confidential owing to two reasons. First, Heppner communicated with a publicly accessible third-party platform. Second, Anthropic's Privacy Policy (in effect at the time of Heppner's use) expressly disclosed that Anthropic collects users' "*inputs*" and Claude's "*outputs*," uses such data to "*train*" Claude and reserves the right to disclose data to "*third parties*," including "*governmental regulatory authorities*." The policy put users on clear notice that Anthropic could "*disclose personal data to third parties in connection with claims, disputes, or litigation" even in the absence of a subpoena. Accordingly, Heppner had no "reasonable expectation of confidentiality*" in his Claude communications having voluntarily disclosed information to the platform which the platform is known to retain in the normal course of its business. The Court also distinguished the AI Documents from a client's privileged notes. Heppner had first shared his "*notes*" with Claude, a third party, before sharing them with his attorneys.

C. Obtaining Legal Advice from Claude

The Court acknowledged that this argument presented "*a closer call*" given the Defendant's assertion that Heppner used Claude for the "*express purpose of talking to counsel*". However, the

Court held that because Heppner acted of his own volition, and not at his lawyer's direction, the relevant inquiry was whether Heppner intended to obtain legal advice from Claude, not whether he later shared Claude's outputs with his lawyers. The Court noted that Claude itself disclaims providing legal advice stating that it is "*not a lawyer and can't provide formal legal advice or recommendations*" and directed users to "*consult with a qualified attorney*". The Court added an important hypothetical, stating that had his lawyers directed Heppner to use Claude, there might have been an argument that Claude functioned as a "*highly trained professional who may act as a lawyer's agent*" within privilege protection. However, that was not the case here.

D. Subsequent Sharing of Information with Attorney does not Confer Privilege

The Court reaffirmed that it is "black-letter law" that non-privileged communications do not become privileged merely by being shared with counsel.

II. Work Product Doctrine

A. Not Prepared by or at the Behest of Attorneys

The Court noted that the doctrine's application to materials in possession of a client depended "upon the existence of a real, rather than speculative, concern that the thought processes of [the client's] counsel in relation to pending or anticipated litigation would be exposed." Even assuming the AI Documents were prepared "in anticipation of litigation," the Court held they were neither "prepared by or at the behest of counsel" nor did they reflect the defence attorney's strategy.

B. Not a Reflection of Lawyers' Mental Strategy

As noted above, the Defendant's lawyers had already conceded that while the AI Documents "*affected*" their strategy going forward, they did not "*reflect*" their strategy at the time Heppner created them. Thus, while the work product doctrine protects the mental processes of the attorney, i.e., the zone of privacy within which a lawyer develops legal theories and strategy; however, those processes were not embodied in documents Heppner generated on his own initiative.

Ramifications

The decision in *Heppner* demonstrates how the use of publicly accessible AI tools can result in loss of privilege and the questions that the use of generative AI will begin to raise in both civil and criminal legal procedure. While the decision will not bind all courts in the United States considering

that it has been delivered by a district court, it still sheds light on several facets that are to be kept in mind while using these tools. Further, it could be of significant persuasive value in Indian courts.

At the outset, the Court's heavy reliance on Anthropic's own Privacy Policy, which permits disclosure to a host of third parties including "*governmental regulatory authorities*" and in connection with "*claims, disputes, or litigation*", introduces a critical practical lesson which showcases that a platform's data-sharing policies can independently defeat claims of confidentiality or privilege. Users of AI tools who input sensitive information are thus presumed to have notice of these policies, regardless of subjective intent to keep communications private.

Additionally, perhaps the most significant aspect of the ruling is the Court's explicit dictum that the outcome might have been different had Heppner's lawyers directed him to use Claude, leaving a narrow pathway open for future litigants in potentially bringing resulting communications within privilege. The ruling carries broad implications for individuals who independently use AI tools to research their legal situation or develop factual narratives. Such individuals now know that those AI outputs could be fair game for government seizure and be used as evidence against them.

Further, under Indian law, the privilege given to advocates is actually narrower in scope than in the United States. Section 29 of the Advocates Act explicitly states that the practice of law in India is exclusively reserved for advocates, and non-attorneys such as accountants and business consultants do not fall under the umbrella of this privilege. Thus, generative AI platforms would fall outside of this privilege if not used directly by or on the instruction of advocates, enhancing the ramifications of this ruling.

Notably, *Heppner* also adds a new leaf to advisory, as clients could be advised not to independently use and provide prejudicial and confidential information to consumer AI platforms to analyse their legal situations as such communications are clearly not protected on many fronts.

This ruling arrives at a moment when AI adoption is accelerating rapidly and, as more individuals use AI tools to navigate their legal issues, such questions of privilege, confidentiality and evidentiary use will multiply. *Heppner* provides the first definitive judicial framework for analysing issues arising out of the independent use of generative AI by a party (who is not an advocate) and is likely to be widely referred to as courts across the world are beginning to encounter similar issues for the first time.

DISCLAIMER

This alert has been written for general information of our clients and should not be treated as a substitute for legal advice. We recommend that you seek proper legal advice prior to taking any action pursuant to this alert. We disclaim all liability for any errors or omissions. For further clarifications, you may write to Pooja Tidke (pooja.tidke@parinamlaw.com) or Rhea Rao (rhea.rao@parinamlaw.com).