akerman

Practice Update

Court Rules That Information Disclosed by Layperson to AI Tools Is Not Protected by Attorney-Client or Work-Product Privileges

February 18, 2026
By Darryl R. Graham

On Tuesday, February 10, 2026, U.S. District Judge Jed S. Rakoff, from the Southern District of New York (Manhattan), ruled that information provided to an AI tool was not privileged and therefore discoverable, even after the client provided the search results to counsel.[1]

In this securities and wire fraud criminal case, the defendant (Heppner) used Anthropic's AI tool, Claude, before his arrest to run queries related to the government's investigation. Critically, Heppner fed information he had learned from his defense counsel at Quinn Emanuel into the AI tool. He then shared 31 AI-generated documents with his defense counsel. When the FBI seized his devices, defense counsel flagged the documents and asserted privilege. The government moved for a ruling that the documents were neither protected by the attorney-client privilege nor the work-product doctrine. After oral argument, the court granted the motion and ordered the disclosure of the 31 documents. In ruling from the bench, Judge Rakoff stated: "I'm not seeing remotely any basis for any claim of attorney-client privilege." In its motion, the government asserted three main arguments.

Related People

Darryl R. Graham

Related Work

Litigation
Securities Litigation

Related Offices

Miami
New York

First, no attorney-client privilege attaches to inquiries made by a layperson to a commercial AI tool, such as Claude (which the defendant used here, and by logical extension, its AI tool peers) because the AI tool is not a licensed attorney. Plainly, there is no attorney-client relationship between an individual and a non-human, commercial AI tool. Further, the consumer-tiered AI tool's terms of service and privacy policy disclaimed any legal advisory relationship, disclaimed any expectation of confidentiality, and explicitly stated that the information provided was subject to use for training purposes and disclosure to governmental authorities.

Second, Heppner's subsequent transmittal of the search prompts and results to counsel cannot retroactively create attorney-client privilege. It is well-settled law that preexisting, non-privileged documents do not become privileged merely because a client later sends them to his or her attorney. The government argued, and the court accepted, that AI prompts and the AI tool's responses should be treated no differently. Rather, the government successfully analogized these AI interactions to conducting a Google search and providing the search results to counsel: no privilege attaches, and the use of an AI tool does not alter this result.

Third, the documents are not protected under the qualified work-product doctrine. On the facts here, counsel did not direct Heppner to conduct the searches in anticipation of litigation. Instead, Heppner independently engaged in these interactions with the AI tool and then later provided the information to counsel. Because the work-product doctrine protects materials prepared "by or at the behest of counsel," Heppner's self-directed AI research falls outside the scope of the work-product doctrine.

Key Takeaways and Practical Advice

This ruling serves as a cautionary tale regarding the use of AI tools and highlights the potential risks arising from their misuse, but, as discussed below, it does not serve as wholesale prohibition.

Here, Heppner, a non-attorney, used what appears to be Anthropic's consumer-tiered version of its AI tool, engaged in independent, self-directed legal analysis, and shared confidential and (what otherwise would have been) privileged information with the AI tool. Based on these facts, the court's core ruling is that these communications and interactions are not privileged and cannot be made privileged by later disclosing them to counsel.

While this area of law is rapidly developing, and this ruling is not binding precedent, it is highly persuasive and, if presented with similar facts, courts are likely to reach the same or similar conclusions in a civil dispute, as the rules governing privilege are the same. Individuals and businesses (including their employees) should be mindful of this when interacting with AI tools and when determining what information is shared with them — especially for consumer-tiered AI tools — because there is legitimate risk that sharing sensitive, confidential, or privileged information may be deemed non-privileged, and therefore discoverable.

Beyond this ruling's direct implication, there are secondary considerations to be mindful of.

What happens if, for example, a client first engages with counsel and receives privileged legal advice, but then shares this advice with an AI tool to understand it better or to inquire about second opinions?

Based on "well-established" legal principles, a client (the holder of the privilege) who voluntarily discloses privileged information to a third party or stranger to the attorney-client relationship (i.e., the AI tool) would likely be deemed to have waived the privilege. On facts similar to *Heppner,* it would not

be surprising if a court were to rule that, had the defendant spoken with his attorney first, but then, independently of counsel, voluntarily shared those privileged communications with an AI tool (especially a consumer-tiered version), the defendant would have likely waived the attorney-client privilege. And, not only would the subsequent exchanges with the AI tool not be privileged, nor would be the underlying advice from counsel, and the waiver could potentially extend to the initial communications between client and counsel that lead to the once-privileged legal advice later shared with the AI tool. In other words, just because a communication is privileged does not mean it will remain so after disclosing it to an AI tool, especially a consumer-tiered AI product.

The terms of service and privacy policy of the AI tools are also important. The terms of service and privacy policy of the popular consumer-tiered AI tools — e.g., OpenAI's ChatGPT and Anthropic's Claude — disclaim any user's expectation of privacy and permit the use of the data provided to the AI tool for data collection, training use, and governmental disclosure. In *Heppner*, the defendant appears to have used a consumer-tiered AI tool, rather than Claude Enterprise, ChatGPT Enterprise, or another API deployment with a zero-data retention policy. These enterprise-tiered AI tools may have materially different terms that could impact the court's analysis on the privilege issue and could, in theory, limit the risk of subsequent disclosure if no data is actually retained by the AI tool. However, the underlying risk regarding privilege — i.e., that no privilege attaches to interactions with an AI tool or that privilege could be waived if disclosed to an AI tool — is not necessarily ameliorated merely by the terms of service of an enterprise product. Therefore, until this precise issue is addressed by a court, there remains risk that such disclosures to and interactions with an enterprise AI tool could be deemed non-privileged and/or result in waiver of privilege.

The solution, however, is not necessarily to avoid AI tools entirely. AI tools are becoming more ubiquitous in business and daily life, and they are objectively helpful in many respects. While clients should exercise caution in the context of consumer-tiered products and should consider implementing internal policies regarding how employees use (or should not use) AI tools, there are ways to use AI tools safely (or at least to mitigate potential risks).

For instance, non-legal business research, internal summarization of public information, and attorney-directed research using enterprise tools with confidentiality provisions would not directly implicate the ruling in *Heppner*. Similarly, the use by in-house counsel of a closed network, enterprise AI tool presents a distinguishable factual and legal scenario that also does not directly implicate *Heppner*. In practice, if a non-attorney intends to engage with an AI tool on legal matters, it is advisable to consult with counsel first, then engage with the AI tool *at the direction of* counsel only (and specifically say so in the prompt with the AI tool), maintain confidentiality with respect to the AI tool's output, and engage only with an enterprise-tiered version of an AI tool (if possible). If an enterprise tool is not an option, it is highly recommended to review the consumer-tiered AI tool's terms of service closely as there are likely limited protections and specific reference to disclosure to governmental authorities. In the latter context, the interactions with the AI tool should be guided by *Heppner*.

This order serves as a reminder that interacting with an AI chatbot, even about legal matters, does not create a privileged communication. Indeed, an AI tool is not an attorney and interactions with an AI tool are not in and of themselves private. Yet, there are ways to mitigate risk and use AI tools in conjunction with counsel to maintain privilege, which then serves to enhance engagement between client and counsel.

As AI becomes further embedded in business and daily life, courts will continue to define the boundaries between technological convenience and longstanding legal protections. Until that line is more clearly drawn, the safest course is to treat AI-assisted legal inquiry as potentially discoverable and to structure its use with that reality in mind.

[1] *United States v. Heppner*, Case No. 25 Cr. 503 (JSR) (S.D.N.Y. Feb. 10, 2026) (docket order granting government's motion at docket entry number 22).

This information is intended to inform firm clients and friends about legal developments, including recent decisions of various courts and administrative bodies. Nothing in this Practice Update should be construed as legal advice or a legal opinion, and readers should not act upon the information contained in this Practice Update without seeking the advice of legal counsel. Prior results do not guarantee a similar outcome.