Proskauer

Recent Federal Privilege Ruling Related to AI Tools Has Implications for Routine Tax Advisor Arrangements

Tax Talks on **February 20, 2026**

On February 10, 2026, Judge Jed Rakoff of the Southern District of New York ruled in *United States v. Heppner* that documents generated through a consumer version of Anthropic's Claude AI were not protected by the attorney-client privilege or the work-product doctrine under the circumstances presented. The decision appears to be the first to squarely address privilege and work product claims arising from a non-lawyer's use of a consumer-grade insecure, non-enterprise AI tool for "legal research," as well as the potential consequences of inputting privileged information (provided to an individual by counsel) into an AI tool. However, putting the novelty of the AI context aside, Judge Rakoff grounded his analysis in traditional privilege principles: that disclosure of privileged communications to a third party in circumstances that undermine confidentiality (here, the corporation operating the AI tool) may result in waiver. And that an AI tool is just that – a tool, not an attorney. Accordingly, this decision reinforces the importance of only using properly secured AI tools with confidential or privileged information and for decisions about using AI in the privileged context to be made by those who best appreciate the risks involved: i.e., lawyers.

What Happened?

After receiving a grand jury subpoena and retaining counsel, the criminal defendant —Heppner — used a non-enterprise, consumer version of Anthropic's Claude to research legal issues related to the government's investigation. Without counsel's direction or involvement, Heppner input information he had learned from his attorneys into the AI tool, generated "reports that outlined defense strategy, that outlined what he might argue with respect to the facts and the law" and later shared those materials with his lawyers. His defense counsel asserted attorney-client privilege and work-product protection for the AI-generated reports, arguing that Heppner had created the AI documents for the purpose of speaking with counsel to obtain legal advice. In response, the government moved for a ruling that the AI documents were protected by neither doctrine, which Judge Rakoff granted.

Key Takeaways

No reasonable expectation of confidentiality.

The court noted that the tool's terms permitted the provider — here Anthropic — to disclose user data to regulators and to use users' prompts and outputs for model training. In other words, the terms themselves made clear that the use of this specific tool was tantamount to a disclosure to the third party that provided the tool. As a result, the court found that users lack a reasonable expectation that their inputs and outputs are confidential. While this reasoning applies broadly to standard consumer AI offerings (which generally provide less confidentiality protections and assurances), the decision leaves open whether enterprise-level products — particularly those that exclude user data from training and provide contractual confidentiality protections — might support a different expectation-of-confidentiality analysis. Importantly, contractual confidentiality protections alone do not automatically establish attorney-client privilege. Even where an enterprise AI product limits data use and includes confidentiality commitments, courts will still assess whether the communication was made for the purpose of obtaining legal advice and whether confidentiality was maintained in a manner sufficient to preserve privilege under governing standards.

Use of unsecured consumer AI tools may defeat privilege.

The court held that discussions with a non-enterprise AI platform are legally equivalent to discussing legal issues with a third party and emphasized how the tool itself disclaimed providing any legal advice. This means that employees using consumer-grade AI tools to analyze legal exposure, assess complaints, research regulatory issues, or prepare for litigation could generate documents that adversaries can later seek to obtain. In this sense, this ruling is consistent with legal ethics opinions that raise the concern that using privileged information with certain unsecured AI tools could be considered a disclosure to the third party that operates those tools, and thus, would be inappropriate for legal work. *See, e.g.*, American Bar Ass'n Standing Comm. on Ethics & Pro. Resp., Formal Op. 512, at 6 (July 29, 2024) ("Self-learning GAI tools […] raise the risk that information relating to one client's representation may be disclosed improperly.").

Lack of attorney direction undermined the work-product claim.

According to the court, because Heppner conducted the AI research independently and not at counsel's direction, the work-product doctrine did not apply. The court indicated that the analysis might differ if the AI use had been directed by counsel under a Kovel-type arrangement: "[h]ad counsel directed Heppner to use Claude, 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." Whether that type of arrangement would result in protection remains an open question.

In an illustration of the anthropomorphism that may happen related to AI tools, the court noted that "what matters for the attorney-client privilege is whether Heppner intended to obtain legal advice from Claude, not whether he later shared Claude's outputs with counsel." [emphasis in original]. Claude, however, is not a person or an attorney. Future cases likely will have to grapple with the question of how consumer AI tools meaningfully differ from other AI tools that are more specifically designed to operate in the legal arena.

Recommended Next Steps

- Be intentional: Reasonable expectations of privacy continue to be of paramount importance when determining whether a tool is suitable for use with confidential or privileged information. Ensure that your organization is conducting proper due diligence when selecting tools and determining permissible applications.

- Audit AI usage policies: Confirm whether your organization permits use of consumer-grade (unsecured) AI tools and make sure that only appropriate applications are allowed – for example, those that do not involve confidential or privileged information.
- Implement guardrails: Restrict input of privileged, confidential, or investigation-related information into consumer AI systems absent a vetted enterprise agreement and clear internal protocols. Given that even the use of a secured AI tool has a risk of privilege waiver depending on the context of the use (for example, inputting a privileged document into a secured AI tool for a non-legal purpose may still present waiver risks depending on the circumstances), require privilege-related decisions to be made by those who best appreciate the risks, such as counsel.
- Train personnel: Ensure employees understand the various considerations that go into determining whether a specific AI tool is appropriate for a specific usage.
- Review Kovel agreements: Kovel agreements among tax advisors and attorneys are widespread, and they continue to be best practice. Evaluate existing agreements to ensure they address use of AI, and update Kovel templates to cover this issue.

What the Decision Does Not Address

The district court's ruling was limited to a criminal defendant's use of a consumer, non-enterprise AI platform without counsel's direction and under terms permitting provider access to user data and does not resolve several important questions.

- As noted above, the decision does not address whether use of an enterprise-tier (i.e., secured) AI product could support a different expectation-of-confidentiality analysis.
- Nor did the court decide whether AI research conducted at the direction of counsel, for example, under a Kovel-type arrangement, or integrated into a structured legal workflow, might qualify for work-product protection.
- The question of whether the same holding would necessarily apply in all civil contexts remains unanswered. The court cited *United States v. Adlman*, 68 F.3d 1495 (2d Cir. 1995), a Second Circuit decision dealing with protection for tax advice given by an accounting firm for a potential corporate merger. Protections for tax advice in the civil context are more robust; for example, Internal Revenue Code section 7525 extends the attorney-client privilege to accountants' tax advice in certain noncriminal tax matters/proceedings, but not criminal ones.
- The opinion also does not establish a categorical rule that all AI-assisted legal work is unprotected; rather, it applies traditional privilege principles to the specific facts before the court.

As AI adoption accelerates, courts are likely to continue scrutinizing how these tools intersect with privilege, confidentiality and waiver doctrines. Organizations should reassess AI governance frameworks now to mitigate litigation and regulatory risk.

The Proskauer team stands ready if you would like assistance reviewing AI usage policies, enterprise agreements, or privilege-protection protocols.

Related Professionals

- **Amanda H. Nussbaum**
  Partner
- **Edward Wang**
  Associate
- **Margaret A. Dale**
  Partner
- **Martin T. Hamilton**
  Partner
- **Nolan M. Goldberg**
  Partner
- **Peter J. Cramer**
  Associate
- **Richard M. Corn**
  Partner