



# Using AI without waiving privilege: Lessons from *United States v. Heppner*

February 19, 2026

Shawn C. Helms | Caitlin Howe | Jason D. Krieser | Joseph B. Evans

## Summary

A recent decision from the US District Court for the Southern District of New York (SDNY), [\*United States v. Heppner\*](#), has generated outsized commentary suggesting that the use of generative AI tools may jeopardize attorney-client privilege. A closer reading shows something far less dramatic. The ruling does not alter black-letter privilege doctrine. Instead, it applies settled principles to a new factual setting: a defendant who voluntarily shared information with a public AI chatbot and later sought to shield the resulting materials from the government.

The lesson is not that AI destroys privilege. It is that longstanding rules about confidentiality, third-party disclosure, and work product apply with equal force when the third party is a chatbot.

## In Depth

### The background

Bradley Heppner, an executive charged in SDNY with securities and wire fraud relating to Beneficient and GWG Holdings Inc., generated a series of documents by querying Anthropic's consumer AI chatbot, Claude. According to reporting, he created 31 AI-generated documents and later shared them with his lawyers.

The government moved to compel production and use of those materials at trial. Heppner argued that the documents reflected legal strategy and were created for the purpose of consulting with counsel and, therefore, were protected either by attorney-client privilege or work-product doctrine.

Judge Jed Rakoff rejected those arguments in a bench ruling on February 10, 2026, holding that the materials were not privileged and not protected work product. Judge Rakoff ruled from the bench, and no written opinion has yet been issued.

### The court's ruling



The decision rests on three familiar pillars of privilege doctrine, as outlined in the following sections.

### 1. No attorney-client communication

Attorney-client privilege protects confidential communications between a client and a lawyer (or their necessary agents) made for the purpose of seeking or providing legal advice.

The AI transcripts at issue were exchanges between Heppner and Claude. Claude is not a lawyer, and it was not engaged by counsel as an agent to assist in providing legal advice. As such, the communications did not fall within the core definition of attorney-client privilege.

The court rejected the idea that a conversation with a commercial AI system could itself qualify as a privileged communication simply because the user intended to later discuss the output with counsel. A client's independent consultation with a third-party information source – whether a search engine, a library, or an AI tool – does not become privileged merely because the resulting notes are shared with a lawyer.

This reflects a basic rule: Privilege protects communications within the attorney-client relationship, not every step a client takes in preparing to speak with counsel.

### 2. Lack of confidentiality and waiver

Confidentiality is central to privilege. **Voluntary disclosure of privileged material to a third party that does not preserve confidentiality generally waives the privilege.**

In *Heppner*, the court focused on the Claude terms of service and privacy policy. In these terms, Anthropic reserves rights to log prompts and outputs, use them for model training, and disclose information to regulators or other third parties. In the court's view, submitting information to a system with express provisions undermining confidentiality was inconsistent with maintaining a reasonable expectation of privacy.

Judge Rakoff reportedly emphasized that the defendant had “disclosed it to a third party, in effect, AI,” under terms that **did not guarantee confidentiality**. That disclosure defeated any claim that the communications were protected.

Again, this is orthodox doctrine. If a client forwards a privileged email to a friend, posts it on a public platform, or shares it with an unprotected third-party service, privilege is typically waived. The fact that the third party here was an AI platform does not alter the analysis.

### 3. No retroactive “cloaking” of privilege



Heppner argued that the AI-generated materials reflected legal strategy and were created for the purpose of seeking legal advice. But the court rejected the notion that forwarding a document to counsel after the fact can retroactively transform a nonprivileged document into a privileged one.

This too is a settled principle. A preexisting document does not become privileged simply because it is later sent to a lawyer. Privilege attaches to communications, not to underlying facts or independently created materials.

#### 4. Work-product doctrine

The court also rejected work-product protection. The key distinction was that Heppner created the materials on his own initiative, not at counsel's direction.

Work-product doctrine protects materials prepared by or for a lawyer in anticipation of litigation. Here, the absence of evidence that counsel directed the AI queries was dispositive. Without that nexus to counsel's litigation strategy, the documents did not qualify.

This aspect of the ruling underscores a practical point: Who directs and controls the preparation of materials matters significantly in work-product analysis.

#### What the case does, and does not, do

*Heppner* does not announce a new rule that communications with AI tools can never be protected. Nor does it suggest that enterprise AI systems, engaged under confidentiality agreements and used under counsel's supervision, are categorically outside the privilege framework.

Instead, the case applies traditional elements of privilege:

- A communication,
- Between privileged persons (or necessary agents),
- For the purpose of legal advice,
- Made with a reasonable expectation of confidentiality.

When those elements are not met, particularly when confidentiality is undermined by voluntary disclosure to a third party, the privilege fails. The technology is new. The doctrine is not.

#### Practice implications



Although the ruling breaks little new legal ground, it has meaningful operational implications.

### 1. Public AI tools pose real waiver risk

Clients and employees may assume that interacting with an AI system is functionally private. Heppner underscores that if the platform's terms permit retention, training use, or third-party disclosure, courts may view the interaction as inconsistent with confidentiality.

From a privilege standpoint, entering privileged communications or legal strategy into a public AI tool can be analogous to disclosing them to an unprotected third party.

### 2. Structure and supervision matter

The work-product analysis highlights the importance of counsel direction. If AI tools are used in connection with litigation or investigations, documenting that such use is at the direction of counsel and within a controlled workflow may be critical. For example, if counsel invites a client into counsel's enterprise version of an AI tool (in the same way users can be invited into a Microsoft SharePoint instance or invited to collaborate on a OneDrive file) behind counsel's firewall, where there is an expectation of confidentiality and training can be turned off in the AI tool, a court may view the traditional privilege arguments very differently.

Courts have long extended privilege and work-product protection to communications involving translators, forensic accountants, consultants, and e-discovery vendors when they are necessary to facilitate legal advice. Whether AI systems can be analogized to such agents will likely depend on certain facts, including contractual protections and supervision.

### 3. Policy and training updates

Companies should:

- Ensure that AI use policies do not allow sharing of confidential and privileged information to public AI models
- Emphasize in employee training that "talking to the AI" is not the same as speaking with legal counsel

## Conclusion

*United States v. Heppner* is best understood as a technology-neutral decision applying longstanding privilege principles to a new context. The court did not rewrite the rules of attorney-client privilege or work product. It enforced them.

For practitioners and clients, the case is a reminder that privilege depends on structure, control, and confidentiality. As



generative AI tools become ubiquitous, preserving privilege will require the same discipline that has always been required when engaging third parties, now applied to the AI world.

## Get In Touch

Shawn C. Helms

[View Profile](#)

Caitlin Howe

[View Profile](#)

Jason D. Krieser

[View Profile](#)

Joseph B. Evans

[View Profile](#)

This material is for general information purposes only and should not be construed as legal advice or any other advice on any specific facts or circumstances. No one should act or refrain from acting based upon any information herein without seeking professional legal advice. McDermott Will & Schulte\* (McDermott) makes no warranties, representations, or claims of any kind concerning the content herein. McDermott and the contributing presenters or authors expressly disclaim all liability to any person in respect of the consequences of anything done or not done in reliance upon the use of contents included herein. \*For a complete list of McDermott entities visit [mwe.com/legalnotices](https://mwe.com/legalnotices).

©2026 McDermott Will & Schulte. All rights reserved. Any use of these materials including reproduction, modification, distribution or republication, without the prior written consent of McDermott is strictly prohibited. This may be considered attorney advertising. Prior results do not guarantee a similar outcome.