

PUBLICATIONS

## Urgent Alert: Federal Judge Rules AI-Generated Documents Are Not Privileged – A Game-Changer for Legal Strategy

February 20, 2026

Geoffrey Vance, Kathryn Johnson, Caroline Boudreau Sweeney



In the fast-paced intersection of artificial intelligence and law, a big decision by a New York federal judge provides a serious warning: carelessly using commercially available generative artificial intelligence tools can destroy the attorney-client privilege and forfeit work product protection, leaving sensitive legal strategies vulnerable to exposure.

### The Case: U.S. v. Heppner – Judge Rakoff’s Rejection of AI Privilege Claims

Here’s what happened. In the case *United States v. Heppner* (S.D.N.Y. 2026), the former CEO of a financial services firm is accused of cheating investors out of \$300 million by lying about a loan tied to a lender he secretly ran. Before he got arrested but after he received a grand jury subpoena, the defendant turned to an AI tool from Anthropic called “Claude.” He typed in details he got from his lawyers and Claude generated 31 documents in response to the defendant’s prompts, including an outline of a defense strategy based on the charges his lawyers anticipated. The Government seized various documents and devices, which included the AI-generated documents, prompting a dispute over whether the documents were protected by the attorney-client privilege and work product doctrine. In a ruling on the bench followed by a February 17, 2026, written opinion, Judge Jed S. Rakoff of the Southern District of New York concluded they were not.

## Attorney-Client Privilege

Judge Rakoff methodically dismantled Heppner's privilege claim, concluding that the 31 AI-generated documents flunked all three elements of the attorney-client privilege:

1. **No Lawyer-Client Communication:** The judge pointed out that Claude is "plainly not an attorney." Privilege only protects discussions with lawyers. The defendant's AI queries were akin to a conversation between two non-attorneys.
2. **No Expectation of Privacy:** The judge pointed to Anthropic's privacy policy, which warns users that the company might share the user's information with others if there's ever a legal fight. He concluded that AI users do not have a reasonable expectation of confidentiality or privacy when communicating voluntarily with a third-party AI platform.
3. **Not for Legal Advice:** The judge noted Claude's explicit disclaimers and warnings against treating outputs as legal advice, urging users to consult their own attorneys. He reasoned that the "predominant purpose" of the defendant's use of Claude could not possibly have been for the purpose of obtaining legal advice when the tool itself expressly warns against such use and says its output should not be considered legal advice.

## Work-Product Doctrine

The defendant's work product argument didn't fare any better. The work product doctrine protects materials that show a lawyer's thinking, made by or for the lawyer. The 31 documents were not prepared by or at the direction of the defendant's attorneys and did not reflect any of their thoughts or strategy. Had the documents been prepared by or at the direction of the defendant's attorneys, the judge would have had to analyze whether disclosure of the materials to Anthropic waived work product protection.

## Why This Matters to You – The High Stakes of AI in Legal Practice

Judge Rakoff's decision isn't just about the defendant, it's a blueprint for how other judges may scrutinize AI in litigation nationwide. It's true that AI is revolutionizing the legal practice from contract analysis to discovery, but the Heppner ruling exposes a glaring vulnerability: **consumer-grade AI can turn strategic brainstorming into discoverable evidence.** Imagine innocently typing case details into Claude, ChatGPT, Grok, or any other of the dozens of publicly available GenAI platforms, only for opponents to seek and receive the output and then use it against the very people who interacted with the platforms.

What's the bottom line, you ask? We asked Claude that exact question. Claude says:

Using a public AI like Claude to process privileged legal information waives attorney-client privilege and eliminates work-product protection, because no lawyer-client relationship exists, confidentiality is destroyed by the platform's privacy policy and data practices, and the communications were not directed by or for obtaining advice from a licensed attorney.

## Actionable Insight

Judge Rakoff's first-of-its-kind ruling teaches that, while AI can be a powerful litigation tool, it is one best left in the hands of attorneys operating in a secure, closed environment that is specifically designed to keep client information confidential and privileged. Otherwise, you may end up prompting AI to write your own smoking gun exhibit.