February 24, 2026

Federal Court Rules Communications with an AI Model About Legal Strategy Are Not Protected by Privilege

By: David A. Bell , Benjamin S. Kingsley , Charles Moulins , Justine Vandermel

What You Need To Know

- A federal court ruled that communications by a non-attorney with a public AI tool about legal strategy are not protected by attorney-client privilege or the work-product doctrine.
- The decision in *U.S. v. Heppner* underscores the risks associated with non-attorneys using AI platforms for legal work, and with giving legal information to AI platforms that lack confidentiality safeguards.

Related Professionals

David A. Bell
Partner · Corp

Benjamin S. K
Partner · Litig

Charles Mouli
Counsel · Litig

- Companies should use secure, enterprise AI tools under attorney supervision to reduce discovery risks and protect privileged information.

In a decision that impacts both attorneys and non-attorneys using artificial intelligence tools to analyze legal matters, Judge Jed S. Rakoff of the U.S. District Court for the Southern District of New York recently ruled in *U.S.* v. *Heppner* that a criminal defendant's communications about legal strategy with Anthropic's large language model (LLM), Claude, were **not** protected against government inspection under attorney-client privilege or the work-product doctrine.

The government sought to examine over 30 communications between criminal defendant Benjamin Heppner, who is charged with securities fraud and related offenses, and the Claude AI platform. Without the involvement of any lawyers, Heppner used Claude to outline his potential legal defense strategy. Heppner claimed that the resulting documents were protected against disclosure to the government because they included information he had learned from counsel, were created to facilitate obtaining legal advice from counsel, and were later shared with counsel.

Judge Rakoff rejected that argument, holding that the government could examine

Related Practices

→ Litigation

Related Industries

→ AI & Machine Learning

Share

→ Subscribe

the communications. Judge Rakoff stated that attorney-client privilege applies only to communications (i) between a client and an attorney, (ii) intended to be confidential, and (iii) made for the purpose of obtaining or providing legal advice. Judge Rakoff then held that Heppner's communications with Claude met none of these conditions. First, he observed that Claude is not an attorney, so by extension Heppner's communications with Claude were not communications between client and attorney. Second, Heppner could not reasonably have expected confidentiality, as Anthropic's then-effective privacy policy made clear that it collected user data and could disclose that data to third parties, including governmental regulatory authorities. Third, the communications were not for the purpose of obtaining legal advice because Heppner did not use Claude at the direction of any lawyer. Judge Rakoff observed that "even if certain information that Heppner input into Claude was privileged, he waived the privilege by sharing that information with Claude and Anthropic."

Separately, Judge Rakoff concluded that Heppner's use of Claude was not protected by the work product doctrine, which he construed to apply only to materials prepared by or at the behest of counsel in anticipation of (or for use in) litigation. Because Heppner's use of Claude was not at counsel's direction, he found it was not subject to work product protection.

Practical Takeaways for Companies Using AI Tools

When Heppner, a non-lawyer, independently developed legal strategy through a public AI model whose terms of service expressly disavowed confidentiality, his communications were discoverable in litigation. Though the law regarding the use of AI tools for traditionally privileged work is unsettled, companies may consider the below measures to help mitigate some of the risks identified by *Heppner* (and to help strengthen their arguments that AI interactions in furtherance of legal work are privileged):

1. For any legal work, use secure, private AI tools with enterprise agreements and policies that safeguard confidentiality.

- Generally, enterprise AI services will be more likely to include confidentiality provisions that safeguard information than general consumer AI products. Companies should therefore consider adopting policies and practices that limit AI use for company business to private, secure, enterprise AI platforms.
- Review the applicable terms of service, privacy policies, and/or enterprise agreement of AI platform providers. Specifically, companies should consider favoring terms that (i) do not allow the sharing of user data with any third parties (including governmental regulatory authorities) or its use in AI model training, and (ii) provide a process for the user to object and assert privilege before information is produced or shared in response to a subpoena.

2. Ensure that lawyers direct and supervise the use of AI for legal purposes.

- Consider adopting policies and practices to prevent the entirely self-directed use of AI by non-lawyers to develop legal strategy. *Heppner* suggests that such uses are unlikely to be privileged.
- Make clear to employees that AI should not be independently consulted for legal advice.

3. Protect the confidentiality of AI prompts and outputs related to legal work.

- Consider limiting data access to lawyers and to employees who reasonably need such information to obtain or facilitate the provision of legal services under the direction of counsel.

4. Ensure that employees do not share any privileged information or communications with a public AI platform, including by educating employees about the attendant risk of privilege waiver.

Unanswered Questions Around AI and Legal Privilege

Many emerging applications of AI in the legal field raise novel challenges that courts will still need to address post-*Heppner*. Courts may continue to look for parallels in existing privilege law to analyze such questions. Companies and their counsel should be attendant to the risks in proceeding in these unresolved areas, which may include:

- Attorneys, or employees at the direction of counsel, using a public AI tool instead

of a private enterprise tool to conduct tasks

- Non-lawyers extensively using AI tools with some initial attorney direction but no meaningful ongoing supervision from lawyers
- Use of AI transcription or translation tools that lack clear confidentiality provisions
- AI platforms whose enterprise agreements include language allowing the provider to disclose user data to government authorities and/or civil litigants in response to subpoenas
  - Agreements that allow users to object to data disclosure in response to such subpoenas may help mitigate that risk
- Internal use of legal-related prompts or outputs to train private AI models

Related Insights

Publications

Publications

Trade Secrets Under Pressure: Lessons from the Coda v. Goodyear Decision

January 27, 2026

White House Executive Order Creates National AI Policy, Overriding States

December 15, 2025

View more related insights