PUBLICATION

Your AI Prompts May Be Discoverable: What Every Client Must Know

Authors: Edward D. Lanquist, Jr, Andrew Jacob Droke, Nicole L. Imhof
March 11, 2026

A recent federal court ruling has delivered a significant wake-up call to anyone who has used an artificial intelligence (AI) platform to research legal issues, understand their rights, or prepare for a conversation with an attorney. In February 2026, United States District Judge Jed S. Rakoff of the Southern District of New York held in *United States v. Heppner* that documents generated using an open and non-enterprise AI platform were neither protected by the attorney-client privilege nor shielded by the work-product doctrine. The implications extend far beyond the criminal case in which the decision arose and affect any individual or business that uses publicly available AI tools to analyze legal exposure, evaluate risk, or prepare for litigation.

In this alert, we identify the legal principles that govern the discoverability of prompts and AI-generated content and offer practical guidance for protecting your legal interests going forward.

For clarity, the key factor in these cases is that the AI systems used were "open" systems in which the prompts and responses may be used to train the large language model (LLM) across the platform. Such a system is very different from a "closed" or "enterprise" system in which the prompts and outputs are used only for the customer's benefit. Courts have generally found that using Outlook or another enterprise email system is sufficiently confidential for communications to be protected. Similarly, courts have generally found that prompts and outputs from a "closed" or "enterprise" system are also confidential.

Background: *United States v. Heppner*

Bradley Heppner, former CEO and board chairman of Beneficient, a publicly traded financial services company, was indicted on October 28, 2025, on charges of securities fraud, wire fraud, conspiracy, making false statements to auditors, and falsification of records in an alleged scheme to defraud investors of approximately $150 million.

After receiving grand jury subpoenas and retaining defense counsel, Heppner – acting on his own initiative and without direction from his attorneys – used Anthropic's Claude, a widely available consumer AI platform, to research legal questions, organize his defense theories, and synthesize information he believed relevant to the government's investigation. Some of the information he provided to Claude reflected conversations he had already had with his lawyers.

When FBI agents arrested Heppner and executed a search warrant at his residence on November 4, 2025, they seized electronic devices containing approximately 31 documents consisting of prompts Heppner entered into Claude and the AI-generated responses he received. Heppner's defense counsel asserted attorney-client privilege and work-product protection over all 31 documents. The government moved for a ruling that the materials were not privileged. Judge Rakoff agreed.

Why the Privilege Claims Failed: Five Critical Lessons

1. No Attorney-Client Relationship with an "Open" AI Tool

As affirmed in *United States v. Mejia*, the attorney-client privilege protects confidential communications between a client and an attorney made for the purpose of obtaining or providing legal advice. An AI platform is not an attorney. It holds no law license, owes no duty of loyalty, cannot form an attorney-client relationship, and is subject to no professional responsibility obligations. As Judge Rakoff held, Heppner's communications were not with his counsel but with a third-party commercial platform – and that threshold defect was fatal to the privilege claim.

Clients must understand that no matter how sophisticated, personalized, or legally focused an AI platform's responses may feel, the interaction is not a privileged communication. The interface may resemble a conversation with a knowledgeable advisor, but it carries none of the legal protections of an actual attorney-client communication.

2. No Confidentiality: Platform Terms Control

Perhaps the most broadly applicable aspect of the *Heppner* ruling is Judge Rakoff's confidentiality analysis. For a communication to be privileged, it must be intended to be, and in fact kept, confidential. Anthropic's privacy policy – publicly available and in effect when Heppner used the platform – expressly states that the company collects user inputs and AI outputs, may use that data to train its models, and reserves the right to disclose such data to governmental authorities and third parties, even in the absence of a subpoena compelling disclosure. Judge Rakoff held that, in light of these terms, Heppner had no reasonable expectation of confidentiality in anything he typed into Claude.

Most consumer AI platforms – including ChatGPT, Gemini, Copilot, and Claude – in their standard (non-enterprise) configurations contain similar data use provisions. Reading the terms of service before using any AI platform to discuss sensitive legal, financial, or business matters is no longer optional; it is essential. Enterprise or API-based configurations often offer more protective terms, but they must be evaluated carefully and confirmed in writing.

3. Not for the Purpose of Obtaining Legal Advice

Even if Heppner had argued that he was essentially using Claude as a legal research assistant, that argument was undermined by the platform itself. Anthropic's own published guidance states that Claude is designed to choose responses that "least give the impression of giving specific legal advice." The tool explicitly disclaims providing legal services.

Judge Rakoff, noting this as a "closer call," nevertheless held that a user cannot claim the privilege for communications made to a tool that expressly disclaims providing the very service the user claims to have been seeking. The government specifically cited Claude's disclaimers in its motion, and the court found them persuasive.

4. No Retroactive Privilege: Pre-Existing Documents Cannot Be Cloaked

Heppner's counsel argued that even if the AI-generated documents were not independently privileged, they became privileged when Heppner later transmitted them to his attorneys. However, Judge Rakoff rejected this argument, consistent with well-settled New York precedent.

Courts have consistently held that sending a pre-existing, unprivileged document to an attorney does not transform it into a privileged communication. The AI-generated documents were created before any attorney reviewed them and were not privileged at the time of creation; transmitting them to counsel afterward could not retroactively supply the missing elements of the privilege.

5. No Work-Product Protection Without Attorney Direction

The work-product doctrine protects materials prepared by or at the direction of counsel in anticipation of litigation. Defense counsel conceded at the hearing that Heppner created the AI-generated documents of his own volition and that his legal team did not direct him to do so. That concession proved fatal to the work-product doctrine claim. Because neither Heppner nor the AI tool is legal counsel, and because Heppner was not working at counsel's direction, the materials did not qualify for work-product protection.

Judge Rakoff further noted that the platform's disclaimer of confidentiality independently undermined any work-product claim. The court left open, but did not decide, whether the analysis might differ if counsel had actually directed the client to conduct AI research.

The Waiver Problem: A Compounding Risk

Beyond the immediate discoverability of the AI-generated materials themselves, the *Heppner* ruling highlights a potentially more serious concern: waiver. Heppner fed into Claude information that he had received from his defense attorneys.

The government argued, and Judge Rakoff agreed, that voluntarily sharing privileged attorney-client communications with a third-party AI platform constitutes a disclosure that may waive the privilege not only over the AI documents themselves but also with respect to the underlying attorney-client communications those documents reflect.

The privilege belongs to the client, but so does the responsibility to protect it. Clients who input information received from their attorneys into a public AI platform risk waiving privilege over those original communications – exposing their entire legal strategy.

This Ruling Extends Beyond Criminal Cases

Although *Heppner* arose in a criminal prosecution, the reasoning applies equally to civil litigation, regulatory investigations, employment disputes, internal corporate investigations, and transactional due diligence.

Any time an employee, officer, director, or individual uses a consumer AI tool to analyze legal exposure, evaluate liability, research employment complaints, or prepare for dispute resolution, they may be creating discoverable records. Opposing counsel in civil litigation, government regulators conducting investigations, and adverse parties in arbitration may seek to obtain AI-generated documents and the prompts used to create them.

Admittedly, Judge Rakoff did not rule on a situation where a client uses a "closed" or "enterprise" system. That issue was not before the Court. If Heppner had used an "enterprise" or "closed" system, the confidentiality arguments would likely have flipped. Uncertainty remains as to whether the fact that the prompts and outputs were not created in response to an attorney's instruction remains an open issue.

Practical Guidance

In light of the *Heppner* decision, we recommend the following steps:

- Do not use consumer AI platforms to research legal questions, analyze your legal exposure, or prepare for conversations with counsel. If you need to organize your thoughts or gather information in anticipation of litigation or a regulatory matter, do so through channels that preserve confidentiality.
- Do not input information received from your attorney into any AI tool. Doing so risks waiving the attorney-client privilege over the underlying communications.
- Read platform terms of service carefully before using any AI tool for business or legal matters. If the platform reserves the right to disclose your inputs to third parties or use them for model training, assume there is no confidentiality.
- If your organization uses enterprise AI solutions, confirm in writing that your contractual terms include robust data confidentiality protections and prohibit use of your data for training or third-party disclosure.
- Consult your attorney before using any AI tool in connection with pending or anticipated litigation, regulatory inquiries, or legal disputes. If counsel directs you to use AI tools as part of case preparation, that direction itself may be important to establishing work-product protection.

For further analysis tailored to your sector and compliance footprint, please contact the authors – Edward D. Lanquist, Andrew J. Droke, Nicole Imhof – or another member of Baker Donelson's AI Team.