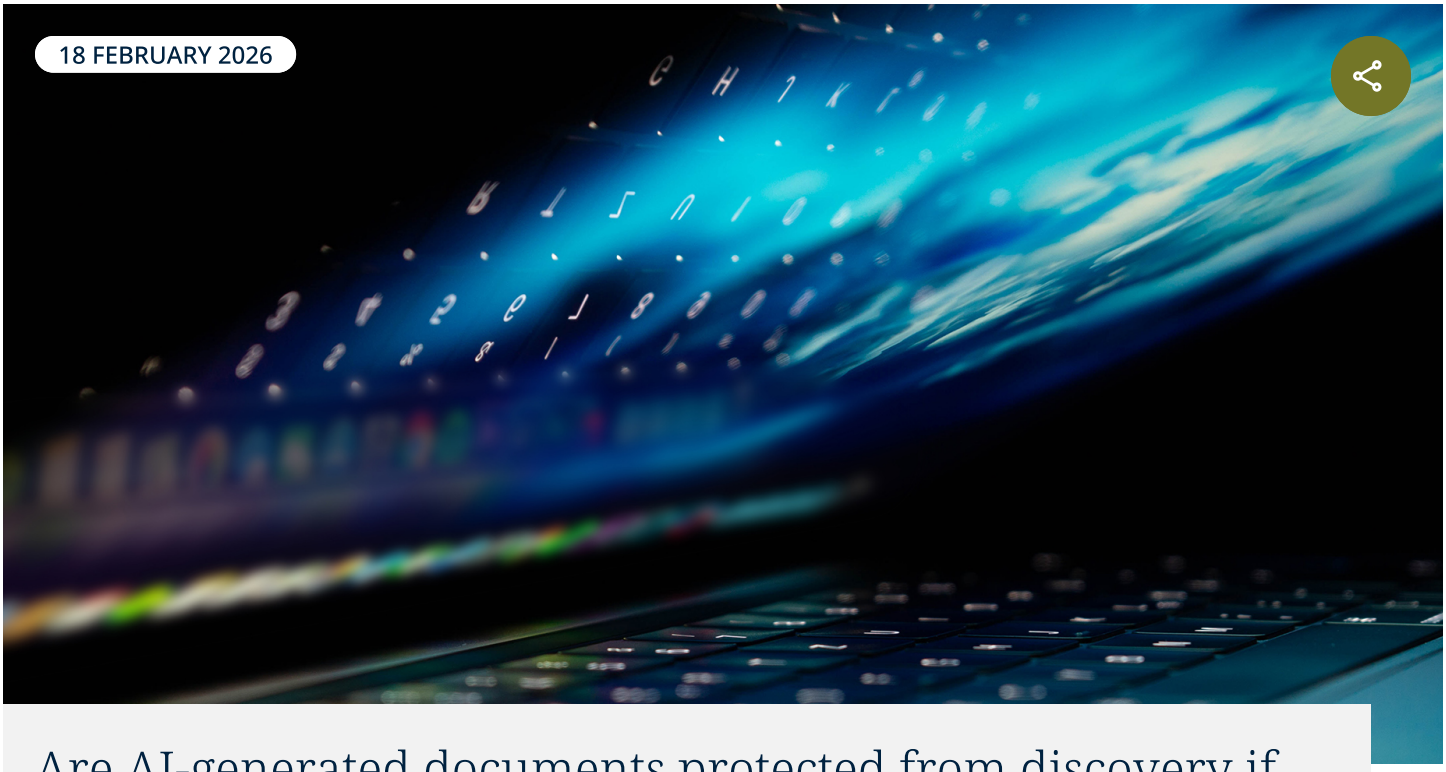

Are AI-generated documents protected from discovery if you send them to your lawyer? One judge says “no”

Key takeaways from the Heppner decision

Written by: Andrew Peck, Danny Tobey, Jason Lewis, Allen Waxman, Michael Atleson, Ryan Lantry, Coran Darling

A federal judge’s recent ruling is a good reminder that, when it comes to discovery issues involving generative artificial intelligence (AI) tools, existing legal standards can sometimes handle them just fine. While some have cast the decision as a problem for privileged AI writ large, the actual ruling stands for a more pedestrian proposition: that documents prepared by a non-lawyer, using a public tool that disclaimed an expectation of privacy, were not privileged.

At a hearing on February 10, 2026, and in a written decision issued on February 17, Judge Jed S. Rakoff of the Southern District of New York considered if the attorney-client privilege or the work product doctrine protects 31 documents that a defendant used a generative AI tool to prepare on his own initiative and then sent to his attorney. The answer? No.

The ruling is notable less for the court’s actual holding than because we still have relatively few judicial applications of attorney-client and work-product principles to the

discoverability of generative AI inputs and outputs. To date, court decisions in this area have stemmed mostly from more specialized situations involving cases against AI companies that own the generative AI tools in question. However, given rapid enterprise adoption and use of generative AI, courts are likely to confront these issues with increasing frequency, including in cases presenting more complex questions. For example, a current area of debate is whether the use of AI tools for litigation-related document review requires the producing party to disclose its AI prompts and outputs.

Nonetheless, as Judge Rakoff stated, the ruling "appears to answer a question of first impression nationwide: whether, when a user communicates with a publicly available AI platform in connection with a pending criminal investigation, are the AI user's communications protected by attorney-client privilege or the work product doctrine?"

Our alert explores the ruling's background and key compliance considerations for businesses and legal practice.

Background

The case involves the criminal fraud prosecution of Bradley Heppner, the former CEO of Beneficient, a Dallas-based financial services company. After learning he was the focus of a law enforcement investigation, Heppner, acting on his own initiative, used a prominent, general-purpose, generative AI chatbot to prepare approximately 31 documents related to his legal case. He later shared these documents with his defense counsel. The government moved for a ruling that these documents are not protected and should be disclosed for trial preparation.

The court's reasoning

Attorney-client privilege does not apply

As Judge Rakoff explained, the attorney-client privilege protects from disclosure "communications (1) between a client and his or her attorney (2) that are intended to be, and in fact were, kept confidential (3) for the purpose of obtaining or providing legal advice."

The court concluded that the documents in question clearly did not satisfy the first two elements because the AI tool was not an attorney and the communications reflected in the documents were not confidential. The lack of confidentiality was apparent from the platform's written privacy policy, which states that it collects user inputs and outputs, uses such data for AI model training, and reserves the right to disclose such data to third parties, including governmental regulatory authorities. Heppner thus had no reasonable expectation of confidentiality.

Judge Rakoff further concluded that the documents failed to satisfy the third element, although he acknowledged that this question "perhaps presents a closer call." Heppner did not communicate with the AI tool "for the purpose of obtaining legal advice." As defense

counsel conceded, while Heppner may have been planning to discuss his findings from the AI tool with his lawyer, he "did not do so at the suggestion or direction of counsel." For privilege purposes, what matters is whether he intended to obtain legal advice from the AI tool, not whether he later shared outputs with his lawyer. As the court stated, "It is black-letter law that non-privileged communications are not somehow alchemically changed into privileged ones upon being shared with counsel."

The court considered, but did not decide, whether the outcome might differ if the defendant's counsel had directed him to use the AI tool, potentially treating the tool akin to a highly trained expert acting as an agent of the attorney, and thus covered by the privilege. Judge Rakoff allowed that this "might arguably" be true, but that these were not the facts present in the case. He also noted that, if prompted for legal advice, the tool will apparently reply that it is not a lawyer and cannot provide such advice, and that users should consult with qualified attorneys instead. (The authors agree wholeheartedly with the AI tool).

The bottom line here, then, is that unprivileged documents that a client creates on his own initiative, even if related to anticipated litigation, do not get stamped with privilege merely because they are later shared with an attorney. The principle applies whether the client creates those documents with an AI tool or crayons.

Work product doctrine does not apply

The defense also argued that the AI documents were protected as work product. Judge Rakoff rejected this argument, stating that the purpose of the work product doctrine is to protect the mental strategies of counsel in anticipation of litigation. Even assuming that Heppner created the documents in anticipation of litigation, defense counsel conceded that the documents did not reflect the attorney's strategy and that the defendant prepared the documents of his own volition.

Implications for legal practice

The *Heppner* decision underscores several key considerations about discoverability when using AI tools in connection with legal matters:

1. **Client communications with AI tools are not privileged.** When a client independently uses an AI tool to analyze legal issues – including issues directly related to pending or anticipated litigation – those communications likely will not be protected by attorney-client privilege. The privilege analysis might be different if the attorney has specifically asked the client to research a particular topic in connection with the representation.

2. **Confidentiality concerns are paramount.** Many public AI tools[1] retain user inputs and outputs, and their terms of service typically permit disclosure to third parties and government authorities. Therefore, users have a diminished privacy

interest in interactions with non-enterprise AI tools. Several state bar associations have cautioned lawyers to use enterprise versions that include protective confidentiality provisions.

3. **Sharing privileged information with AI companies may waive privilege.** If a client inputs information learned from privileged communications with counsel into a public AI tool, the client may waive any applicable privilege by sharing it with the AI company, which is a third party. In a footnote, Judge Rakoff stated that such a waiver would have occurred if Heppner had input information that his attorneys had conveyed to him. This risk is minimized, if not necessarily removed, when using enterprise AI tools that have controls in place to prevent sharing with parties outside of the privileged relationship.

4. **Attorney direction can matter for work product.** Had defense counsel directed Heppner to run the searches, the work product analysis, and potentially the privilege analysis, might be different. For both analyses, the distinction between public and enterprise AI tools may be significant.

Key considerations for businesses

The outcome in *Heppner* highlights key considerations for companies whose non-legal staff use generative AI tools in connection with anticipated or pending litigation:

1. **Assume no privilege with public AI tools.** A company's independent use of public AI tools to research legal matters may not be privileged. Providers of those tools are third parties, and they may retain inputs and potentially disclose them to other third parties.

2. **Exercise caution with confidential information.** Do not input confidential information into public AI tools without understanding the potential for its later disclosure.

3. **Leverage enterprise and, as appropriate, law-specific AI tools where possible.** Enterprise AI tools often offer technical measures and contractual obligations that segregate input and output data from third parties and enable companies to maintain privilege and confidentiality of information. Companies are encouraged to carefully review vendor terms and infrastructure to ensure suitable degrees of segregation.

4. **Document attorney direction.** If appropriate AI tools are to be used as part of litigation preparation, ensure that the use is at counsel's specific direction and properly documented, which may help support a claim for protection.

5. **Segregate certain AI-generated materials.** Treat documents generated by

public AI tools as potentially discoverable and consider their use accordingly. Consider controlled user access to law-specific AI tools, as well as training for non-lawyers on responsible use of AI, including how to protect legal privilege.

Judge Rakoff's concluding paragraph is instructive and succinct, so we use part of it here as well:

Generative artificial intelligence presents a new frontier in the ongoing dialogue between technology and the law. Time will tell whether, as in the case of other technological advances, generative artificial intelligence will fulfill its promise to revolutionize the way we process information. But AI's novelty does not mean that its use is not subject to longstanding legal principles, such as those governing the attorney-client privilege and the work product doctrine.

Learn more

For more information, please contact the authors or visit our Artificial Intelligence and Data Analytics capability page.

[1] A public AI tool is a consumer-facing tool generally made available on a commercial AI platform, often for free or via low-cost subscriptions. By contrast, enterprise AI tools are designed for commercial use by organizations and typically offer enhanced security, data isolation, dedicated infrastructure, and terms that contractually prohibit user data from being combined, used by third parties, or used to train and improve the services. Enterprise AI tools also often contain terms that provide companies with ownership rights to all outputs.

Related insights

FTC resolves another case involving "AI-washing"

5 FEBRUARY 2026

How can in-house lawyers use Generative AI responsibly and effectively?

4 FEBRUARY 2026

Critical audit of NYC’s AI hiring law signals increased risk for employers

30 JANUARY 2026

Related capabilities

- Litigation, Arbitration and Investigations
- AI Litigation
- Technology
- Artificial Intelligence and Data Analytics
- Product Liability
- Business and Commercial Litigation
- Intellectual Property
- Data, Privacy and Cybersecurity

People

Capabilities

About us

Insights

Careers

Locations

News

Events

Blogs

Alumni

Pro bono





Contact us

Find an office

Subscribe

Also of interest

News

Events

Capabilities

Make a payment

Sitemap