WILENTZ
—ATTORNEYS AT LAW—
WILENTZ, GOLDMAN & SPITZER, P.A.

Is Your AI Chat History Discoverable? United States v. Heppner Says Yes

Monday, August 24, 2026

Employers and employees increasingly type questions into consumer artificial intelligence ("AI") chatbots as much as they once used search engines, including questions about workplace disputes, terminations, and legal exposure. A recent decision from the Southern District of New York confirms that *those prompts and the AI's responses are ordinary evidence*. In *United States v. Heppner*, No. 25-cr-00503-JSR (S.D.N.Y. Feb. 10, 2026), a federal judge held that documents generated using an open, consumer version of an AI platform were not protected by the attorney-client privilege nor the work-product doctrine, even though some of the prompts contained information the client had learned from his own lawyers.

The Facts Behind the Heppner Decision

Bradley Heppner, the former CEO of a publicly traded financial services company, was indicted in October 2025 on securities fraud, wire fraud, and related charges. After receiving grand jury subpoenas and retaining defense counsel, Heppner, on his own initiative and without direction from his attorneys, used the consumer version of Anthropic's AI "Claude" to research legal questions, organize defense theories, and synthesize information about the government's investigation. Some of what he typed or uploaded reflected conversations he had already had with counsel. When agents executed a search warrant at his home, they seized devices containing roughly 31 documents of Heppner's prompts and the AI-generated responses. His lawyers asserted privilege and work-product protection over all of them. The government moved for a ruling that the materials were fair game, and the judge agreed.

Why the Court Rejected Privilege and Work-Product Protection

The court's reasoning applied familiar doctrine to a new tool. First, a chatbot is not an attorney because it holds no license, owes no fiduciary duty, and cannot form the trusting professional relationship every recognized privilege requires. Second, there was no confidentiality. The platform's publicly available consumer terms disclosed that user inputs and outputs may be collected, used to train the model, and disclosed to third parties, including government authorities, thus Heppner had no reasonable expectation that anything he typed would stay private. Third, the platform itself disclaims providing legal advice, undercutting any claim that the exchanges were made to obtain legal advice. Fourth, forwarding the documents to his lawyers afterward changed nothing, as sending a pre-existing, unprivileged document to counsel does not retroactively make it privileged. Since Heppner's counsel conceded they never directed him to use the tool, the work-product doctrine could not apply. Even if they had directed him to use the tool, the other concerns raised above still exist.

The court also flagged that by feeding privileged attorney communications into a third-party platform, a client may waive the privilege over the underlying communications themselves, not just the AI transcripts. Notably, the ruling addressed an "open" consumer system whose terms permit training and disclosure. The court did not decide how a "closed" enterprise deployment, where prompts and outputs are used only for the

customer's benefit, would fare, and the confidentiality analysis could well come out differently for such a system.

What the Heppner Decision Means for Employers

Heppner is a criminal case, but its logic governs civil litigation, arbitrations, and agency investigations alike. Courts are already treating generative AI data such as prompts, outputs, and usage logs as electronically stored information subject to the usual relevance and proportionality rules, and have compelled production of AI logs on a massive scale in other litigation. Picture a supervisor asking a consumer chatbot how to "document a problem employee before firing her," or an HR manager pasting a draft investigation summary into a free AI tool. In the discrimination suit that follows, those exchanges are discoverable, unprivileged, and quotable.

Key Employer Takeaways

Employers should treat AI use as both a privilege problem and a discovery problem. Adopt and enforce a policy directing that sensitive legal, personnel, and investigative matters be handled only through counsel or through closed, enterprise AI tools with written confidentiality and no-training commitments, never through personal or free consumer accounts. Inform managers that a chatbot conversation is not a private brainstorm. When litigation is anticipated, extend litigation holds to AI prompts, outputs, and activity logs, coordinating with IT on where that data lives and how long the platforms retain it.

The Employment Law Team at Wilentz is available to assist with the development of workplace AI policies and the management of discovery obligations involving AI-generated content. Call us today to discuss your workplace AI policies.

Attorney

- Tracy Armstrong

Practice

- Employment Law