

# Beware: Use of Generative AI Tools for Sensitive Information is Risky Business

Jeffrey I. Snyder, Kenneth Duvall & Ethan D. Schwab

For many, generative AI tools have augmented, if not replaced, aspects of personal and corporate workflow. As of early 2026, between 37 and 50 percent of people use an AI tool as their primary source of information.<sup>1</sup> In many industries and daily practices, using AI to brainstorm, streamline communication, and solve problems is not only permitted, but expected. Recognizing that free, open-source generative AI tools (i.e., ChatGPT) use submitted data to train the model for future responses, many businesses have developed or subscribed to enterprise-tier, closed-end, proprietary AI tools. These prompts and responses are often logged and retained, thereby creating a record that can later be accessed and potentially discovered.

So it follows that a user's prompts and responses are becoming the subject of questioning and depositions, requests for production, and other discovery requests. Is this information discoverable? Is it admissible? And, if someone crafts a prompt or otherwise uploads sensitive information to a generative AI tool in search of legal advice, is that information protected by the attorney-client privilege?

## Artificial Intelligence and the Attorney-Client Privilege

As a general proposition, where a business user interacts with a generative AI tool (whether or not closed-end) other than possibly at the direction of a licensed attorney, that information is likely discoverable to the same extent that it would be if the same interaction occurred by e-mail between two non-lawyers. The attorney-client privilege protects communications between a lawyer and a client regarding legal advice and, where the communication does not involve a lawyer or legal advice, its discoverability is plainly not precluded by the attorney-client privilege.<sup>2</sup>

Can discovery be avoided when a business user appears to seek legal advice from a generative AI tool? Generally, three elements must be met for a communication to be considered privileged and therefore withheld from discovery: (1) the communication must be between an attorney and a client; (2) the communication must be intended to be, and actually is, kept confidential; and (3) the purpose of the communication is to obtain or provide legal advice.<sup>3</sup>

Though AI tools have developed expertise in many professional fields, and despite people treating these tools as personal, trusted advisors—they are not attorneys. In a recent case out of the federal trial court in Manhattan, *U.S. v. Heppner*, a criminal defendant who anticipated being accused of various offenses in connection with an over \$100 million corporate fraud had conversations with a publicly-available consumer version of an AI chatbot, outlining what would constitute his defense strategy and the potential charges that may be levied against him. In a memorandum explaining its ruling in *U.S. v. Heppner*, Judge Jed S. Rakoff of the United States

District Court for the Southern District of New York concluded that “[b]ecause Claude is not an attorney ... that alone disposes of Heppner’s claim of privilege.”<sup>4</sup>

Whether Claude is a licensed attorney, however, is only the tip of the iceberg with respect to a privilege analysis. What occurs when a communication is made between a lawyer and a client, but the client subsequently uploads that communication to generative AI? The *Heppner* court noted that the confidentiality element was breached by uploading the information to the generative AI chatbot, reasoning that many large language models learn from and train on user-inputted data, and that one does not enjoy a reasonable expectation of privacy in information uploaded to these tools.<sup>5</sup> Although this may appear to establish a blanket waiver of privilege for any information uploaded to generative AI, emerging private AI models offer businesses access to large language models within a company’s secure environment or servers, blurring the line as to whether the information has been shared with a third party.

The third element of privilege—whether the communication is made with the purpose of seeking legal advice—is similarly challenging to apply to AI tools. While the *Heppner* court noted that this element was not met because the defendant’s communications with the AI were not made at the direction of counsel, the court acknowledged that it would have been a closer call had the attorney instructed such chats to occur because the AI chatbot could arguably have been acting as a lawyer’s agent.<sup>6</sup>

The *Heppner* court observed that the defendant having subsequently shared the information generated by the AI chatbot with his counsel did not retroactively “cloak” the information in privilege (although presumably the additive portions of the interaction between the lawyer and client would have been privileged).<sup>7</sup>

### **Artificial Intelligence and the Work Product Doctrine**

Similar to the attorney-client privilege, the work product doctrine protects the mental impressions and thought processes of an attorney, allowing counsel to prepare the client’s case without such impressions being subject to discovery by the opposing party.<sup>8</sup> The *Heppner* court concluded that the defendant’s chats were not protected work product because, even if they were prepared in anticipation of litigation as is required for protection, they were not prepared by or at the direction of counsel.<sup>9</sup>

But a recent case from another federal court—this one out of *Michigan, Warner v. Gilbarco*—appears to take a different approach. In *Warner*, the court held that documents uploaded to ChatGPT were protected by work product and that the protection was not waived by the upload—reasoning that waiver of work product protection requires disclosure to an *adversary*.<sup>10</sup> ChatGPT, the *Warner* court reasoned, “[is a] tool[ ], not [a] person[ ]” and the defendants’ theory that AI use waives work product “would nullify work-product protection in nearly every modern drafting environment, a result no court has endorsed.”<sup>11</sup> These divergent outcomes underscore that both the legal framework and the assumptions informing judicial analysis of these tools remain unsettled and evolving.

### **Bilzin Sumberg’s Confidential, Attorney-Directed AI Use May Preserve Privilege**

Where appropriate, Bilzin Sumberg uses enterprise-grade AI platforms that, among other things, prohibit the retention and disclosure of user data to third parties, and only uses these platforms at the direction of a licensed (human) attorney. Accordingly, the resulting interactions (i.e., prompts and responses) are more likely to satisfy each of the elements of the privilege or work

product doctrines or, if they fail to do so, it will not be because of the manner in which the AI platform was utilized or operates (as opposed to, for example, the potential failure of the work product doctrine to apply in a transactional matter where litigation was not reasonably anticipated).

## **Conclusion**

Until courts establish clearer guidance regarding both the attorney-client privilege and the work product doctrine, businesses and individuals should exercise extreme caution before uploading confidential or sensitive information to any generative AI platform and should consult with legal counsel regarding such use. Clients may also want to implement internal policies governing when and how AI may be used for legal work, including, for example, guidance on platform selection, data handling, and the role of in-house or external legal counsel since courts look at the nature, timing, and objective of an attorney's involvement, not merely than an attorney was involved. When seeking legal advice or preparing for potential litigation, communicating directly with retained legal counsel and seeking legal advice only from a licensed attorney is the safest course of action to preserve the protections that the attorney-client privilege and the work product doctrine are designed to provide.

## **Related Practices:**

[Technology](#), [Data Privacy & Security](#), [Commercial Litigation](#)

## **Related People**



[Jeffrey I. Snyder](#)

[Partner, Bankruptcy & Restructuring, Trial & Litigation, CMBS](#)



[Kenneth Duvall](#)

[Partner, Trial & Litigation](#)



Ethan D. Schwab

Associate, Trial & Litigation

<sup>1</sup> <https://searchengineland.com/consumers-start-searches-ai-not-google-study-467159>.

<sup>2</sup> [https://www.law.cornell.edu/wex/attorney-client\\_privilege](https://www.law.cornell.edu/wex/attorney-client_privilege).

<sup>3</sup> *United States v. Mejia*, 655 F.3d 126, 132 (2d Cir. 2011).

<sup>4</sup> *U.S. v. Heppner*, --- F.Supp. 3d. ----, 2026 WL 436479 (S.D.N.Y. Feb. 17, 2026).

<sup>5</sup> *Id.* at \*6.

<sup>6</sup> *Heppner* at 7.

<sup>7</sup> *Heppner* at \_\_\_.

<sup>8</sup> [https://www.law.cornell.edu/wex/attorney\\_work\\_product\\_privilege](https://www.law.cornell.edu/wex/attorney_work_product_privilege).

<sup>9</sup> *Heppner* at 9.

<sup>10</sup> *Warner v. Gilbarco*, 2026 WL 373043, at \*4 (E.D. Mich Feb. 10, 2026) (*emphasis added*).

<sup>11</sup> *Id.*