ALM | LAW.COM

New York Law Journal

'United States v. Heppner': Generative AI and its Pitfalls for the Attorney Client Privilege and Work Product Doctrine

By Matthew Dunn
April 9, 2026

Introduction

Generative AI tools and platforms are being used in many ways in all facets of business and daily lives. Publicly available AI tools, such as ChatGPT, Claude, and Gemini, are used to answer questions, draft documents, summarize data, among other uses. These public facing generative AI tools use technology—such as large learning models (LLMs) – to train and learn from the data imported by users. While there exist a variety of types of AI platforms with varying security features, users must understand the associated risks.

A recent federal court decision—*United States v. Heppner*—underscores a critical risk: user inputs into public generative AI platforms and output documents derived from those interactions are not necessarily confidential and, in the context of litigation, may not be protected from disclosure by the attorney-client privilege or the work product doctrine.

Summary of the Case and Holding

In *United States v. Heppner*, Judge Jed Rakoff of the U.S. District Court for the Southern District of New York addressed whether documents reflecting a client's interactions with a generative AI platform—later provided to defense counsel—were protected from disclosure to the government (the adversary in this criminal proceeding) by the attorney-client privilege or work product doctrine.

Taking Gen AI Projects One Step at a Time

In this criminal law proceeding, the defendant Bradley Heppner was charged with securities fraud and related crimes. After the defendant

had received a grand jury subpoena, but before he had been indicted by a grand jury (charging him with the crimes), he used a publicly available generative AI tool, Claude (owned and operated by Anthropic), to ask questions and generate documents relating to possible legal defenses and strategies in connection with the criminal charges that might be brought. Importantly, this was not done at the direction of, or in consultation with, Heppner's lawyers. He then later shared those AI-generated documents with his attorneys.

Following his indictment, in executing a search warrant at Heppner's home, the FBI seized documents and electronic devices, which included the inputs (communications) with the AI tool as well as the outputs, all of which related to legal strategies and defenses.

The defendant, through counsel, asserted both attorney-client privilege and work product protection over such documents to prevent their use by the government, arguing that: "(1) Heppner had inputted into Claude, among other things, information that Heppner had learned from counsel; (2) Heppner had created the AI documents for the purpose of speaking with counsel to obtain legal advice; and (3) Heppner had subsequently shared the contents of the AI documents with counsel."

Attorney Client Privilege

· As stated by the court, the attorney-client privilege is a well-established legal protection against disclosure of, "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 rejected the application of the privilege here because the communications were not made to an attorney or an agent of an attorney for the purpose of obtaining legal advice. The AI platform was a third-party service, not a privileged intermediary.

· In addition, the court held that the documents were not confidential because they were disclosed to a third-party AI platform, which has a privacy policy stating that users consent to Anthropic's use of the data (inputs and outputs) to train Claude and reserves the right to disclose data to certain third parties. Any expectation of confidentiality was effectively waived.

· In addition, the court emphasized that Heppner did not create these documents "at the suggestion or direction of counsel" and it did not matter that they were later shared with counsel.

Work Product Doctrine

· The court also summarized the equally well-established work product doctrine, which protects against disclosure of materials "prepared by or at the behest of counsel in anticipation of litigation or for trial." It is intended to protect the mental processes of the attorney but can protect client documents prepared at counsel's direction or request.

· The court rejected application of the work product protection because the documents (inputs and outputs) were not prepared by or at the direction of counsel and did not reflect counsel's strategy. Instead, they were created independently by the client on his own volition. The court did not need to determine if the documents were in fact created in anticipation of litigation.

Takeaways and Lessons

· The decision reinforces a logical principle: once a party voluntarily exposes information to non-confidential third party, privilege protections

are unlikely to be preserved. This can result in critical consequences in litigation. AI tools may be treated as third parties.

· In the use of generative AI tools and platforms, users must be aware of the terms of use, privacy policies, and security features. A key aspect of the court's reasoning—and of broader risk management—is the distinction between open AI platforms and closed, secure AI platforms. Users of AI must understand the types of platforms they are using.

Open AI Platforms, such as internet-accessed Claude used by Heppner, use inputs and outputs to train the systems and they may be shared with third parties. Use of such publicly available tools is not truly confidential. The terms of use and privacy policies generally make these disclosures to users.

Closed or Secure AI Platforms, including product offerings from the popular generative AI tools and other companies, provide secure and controlled environments in which there may be a greater expectation of privacy and confidentiality with respect to inputs and outputs.

These platforms generally are not used for training of data, access is restricted (password protected), and they have security protections and controls. These tools, which are generally more expensive, should be considered for use involving confidential or privileged information. However, it is important to understand the terms and security features before using any such tool.

· Clients should retain and consult counsel before using AI to analyze legal issues or explore legal strategy in connection with an ongoing, threatened, or reasonably anticipated litigation.

· While the *United States v. Heppner* decision was in the context of a criminal proceeding involving an individual, these lessons can be equally applicable to businesses and organizations (including non-profits) and in the context of civil litigations.

· Organizations should conduct training for and provide guidance to employees on AI use. Organizations should consider implementing and disseminating AI use policies applicable to employee use of AI.

· Lawyers must also be aware of AI tools they or their staff may use. Lawyers have ethical obligations to maintain client confidentiality, including confidentiality of communications and documents. Lawyers also have an ethical obligation of technological competence, which requires lawyers to keep abreast of changes in the law and its practice, including the benefits and risks of relevant technology.

Matthew Dunn *is a litigation partner and Chair of Cybersecurity & Data Privacy Practice at Carter Ledyard & Milburn. He serves in an advisory role, counseling clients on security issues and compliance with always-evolving regulatory frameworks, including U.S. and international laws and regulations relating to data privacy, cybersecurity, and artificial intelligence.*