Federal Court Rules Client's AI-Generated Materials Are Not Protected by Attorney-Client Privilege or Work Product Doctrine

Lowenstein Sandler Client Alert

Are client queries with publicly available AI platforms shielded by the attorney-client privilege or the work product doctrine? No, said the first court to consider this issue.

Background

In *United States v. Heppner*,[1] a criminal matter, District Judge Jed S. Rakoff ruled that defendant Benjamin Heppner's communications with Claude, private company Anthropic's AI platform, were not protected by either the attorney-client privilege or the work product doctrine. After Heppner learned that he was the target of a federal securities and wire fraud investigation, he used a public version of Claude to prepare reports outlining possible defense strategies, including what he planned to argue with regard to the facts and law. Heppner used Claude of his own volition, without direction from his lawyers.

When the FBI executed a search warrant at Heppner's home, the agents seized 31 documents that memorialized his communications with Claude. Heppner argued the documents were protected by the attorney-client privilege and the work product doctrine because he had entered information he learned from his lawyers into Claude, created the documents for the purpose of obtaining legal advice from his lawyers, and later shared the contents of the documents with his lawyers.

The court rejected each of these arguments and concluded the documents were discoverable.

Attorney-Client Privilege

The attorney-client privilege protects communications between a client and their attorney that are intended to be and in fact are kept confidential for the purpose of obtaining or providing legal advice.

The *Heppner* court reasoned that because Claude is not a lawyer, Heppner's communications with it were not confidential. Moreover, because Heppner was not seeking legal advice from a lawyer, the 31 documents memorializing his communications with Claude were not privileged. In fact, Claude's written privacy policy–to which Heppner agreed before using the AI tool–specifies that Anthropic collects data on users' inputs and Claude's outputs, uses that data to train Claude, and can disclose that data to third parties, including government authorities. Heppner's sharing the documents with his defense counsel after their creation did not make the unprivileged documents privileged.

Work Product Doctrine

The work product doctrine protects from disclosure materials prepared in anticipation of litigation by a lawyer or at a lawyer's direction. The *Heppner* court explained that the AI-generated documents were not protected by the work product doctrine because Heppner used Claude of his own accord, not at his lawyers' direction or as his counsel's agent, and the documents did not reflect his lawyers' legal strategy.

While the *Heppner* case was a criminal matter, it is not unreasonable to expect the same logic could apply to civil litigation–that is, that one's AI searches could be discoverable.

Takeaways

Consult With Counsel *Before* Using AI: If you are engaged in criminal or civil litigation, are under investigation, or believe you may be charged, you should consult your lawyer before using any AI platform to research or analyze the legal issues and facts of your case.

AI Privacy Policies May Defeat Privilege: The privacy policies of publicly available AI platforms typically permit the collection, retention, and disclosure of user inputs and AI tool outputs to third parties, including governmental authorities. Sharing privileged information with a public AI platform or a company's enterprise AI platform may waive privilege, just like third-party disclosures to humans do. To that end, an employee using their company's enterprise version of an AI platform could be construed as sharing their AI prompts and searches with the entire company.

Consumer AI vs. Enterprise AI: Many enterprise AI platforms contractually guarantee that user inputs will not be used for training and will be kept confidential. Certain enterprise agreements require data isolation and prohibit disclosure, helping maintain confidentiality that publicly available AI platforms may not offer. In its opinion, the court did not address the use of enterprise AI platforms. But an employee should not have an expectation of privacy in their AI searches. Theoretically, the employee's company could later access and review those searches.

Update AI Usage Policies and Training: Companies should work with legal counsel to update their policies governing employee use of generative AI tools so that they specifically address the risks highlighted by the *Heppner* court's decision. Training programs should emphasize that employees should not enter information related to legal matters, investigations, pending litigation, or communications with counsel into external AI platforms and that even in an enterprise platform, entering such information carries risk of subsequent disclosure.

The Lowenstein Sandler Data Privacy, Security, Safety & Risk Management, White Collar Defense, and Employment groups would be pleased to assist in reviewing your policies or answer any questions you may have about this matter.

---

[1] No. 25-cr-503 (S.D.N.Y. Feb. 17, 2026).