

Legal Updates

Court Rejects Privilege Claim Over AI-Generated Documents

February 17, 2026

A recent federal court ruling held that AI-generated documents prepared by a defendant and later shared with legal counsel were not protected by attorney-client privilege or the work product doctrine. In *USA v. Heppner*, Judge Jed S. Rakoff of the US District Court for the Southern District of New York ordered the production of 31 documents generated by the defendant using Anthropic's Claude, notwithstanding that the defendant had incorporated information obtained from counsel and later shared the AI outputs with his attorneys. The ruling—issued from the bench—provides early judicial guidance on privilege issues related to AI-generated materials, an area with limited existing case law.

Background

Bradley Heppner, the former CEO of Beneficient, is charged with defrauding investors and misappropriating more than \$300 million from GWG Holdings, a public company that invested heavily in Beneficient. Following Heppner's arrest, authorities seized electronic devices containing the AI-generated documents at issue. Defense counsel informed prosecutors that the documents consisted of prompts and outputs from the Claude AI tool and asserted that they were privileged because they incorporated information learned from counsel and were intended to facilitate the provision of legal advice. Counsel conceded that they did not direct Heppner to use the AI tool and were not involved in creation of the documents.

Attorney Client Privilege

The court agreed with the government that the AI-generated documents were not protected by the attorney-client privilege, emphasizing three principal considerations:

1. **No attorney involvement:** Communications with an AI tool are not communications with counsel. The court noted that the AI provider's terms of service expressly disclaim any

attorney–client relationship and state that the tool does not provide legal advice.

2. **Not generated to obtain legal advice from a lawyer:** Although Heppner generated the documents after receiving a grand jury subpoena and engaging counsel—and later shared those documents with his attorneys—the prompts and outputs themselves were not communications with counsel. The government emphasized that legal advice cannot be obtained from an AI tool, while the defense argued that the documents merely consolidated Heppner’s thoughts for later discussion with counsel wherein counsel’s legal advice would be sought.
3. **Lack of confidentiality:** The court found that the AI tool’s privacy policy permitted disclosure of user inputs and outputs to third parties, including government authorities, defeating any reasonable expectation of confidentiality. To the extent Heppner input information originally obtained from counsel, the government argued that any privilege was waived when that information was disclosed to a third party. The government analogized Heppner’s use of the AI tool to conducting internet research or consulting library materials.

Judge Rakoff concluded he could not find any basis for a claim of attorney–client privilege over the AI-generated materials and that merely transmitting documents to counsel after their creation did not convert them into privileged communications.

Work Product Privilege

The court also rejected claims of work product protection. Relying on *In re Grand Jury Subpoenas Dated Mar. 19, 2002 and Aug. 2, 2002*, 318 F.3d 379, 384 (2d Cir. 2003), the government argued that the work product doctrine applies only to materials prepared “by or at the behest of counsel in anticipation of litigation.” Here, Heppner independently generated the documents without counsel’s involvement. The court held that independent research, even if later shared with an attorney, does not qualify as protected work product.

It is unclear to what extent defense counsel identified legal support for the proposition that the work product doctrine can protect documents generated by non–attorney parties themselves. *See, e.g., Wultz v. Bank of China Ltd.*, 304 F.R.D. 384 (S.D.N.Y. 2015) (holding that Federal Rule of Civil Procedure 26(b)(3)(A) accords the protection to material prepared “by or for [a] party or its representative”—not merely material prepared by or for an attorney, and citing many cases that support the Advisory Committee Notes clarification that the Rule was intended to also protect material prepared by non–attorneys). Other caselaw exists in this court to affirm that a criminal defendant’s own documented recollection of events, absent any reflection of attorney strategy or thought process, is protectable work product, even more so in the criminal context. *See US v. Stewart*, 287 F. Supp. 2d 461, 465–466 (S.D.N.Y. Oct. 20, 2003) (relying on the language of Federal Rule 26(b)(3) to protect a document prepared “by...a

party”).

Practical Guidance and Best Practices

The decision offers early, practical guidance for lawyers and clients evaluating privilege claims over AI-generated content:

- **Preserve confidentiality:** Use closed-model AI tools with strict privacy policies that prohibit disclosure to third parties or government authorities and do not use any of the user inputs to train models.
- **Ensure counsel direction and involvement:** To support a privilege claim, AI use should occur at the express direction of counsel and for the purpose of facilitating legal advice. Where appropriate, prompts may reflect that the work is being performed at counsel’s direction.
- **Limit dissemination:** Restrict access to AI prompts and outputs to preserve confidentiality.
- **Draft privilege logs carefully:** Privilege logs should clearly articulate the asserted basis for protection, including establishing the criteria for confidentiality safeguards and the role of counsel.
- **Proceed with caution:** Until the law develops further, attorneys should advise clients that inputting otherwise privileged information into AI tools may risk disclosure.

Looking Ahead

Courts are likely to confront privilege issues relating to AI-generated materials with increasing frequency. Because privilege determinations are jurisdiction-specific and highly fact-dependent, future decisions will further define the contours of protection, but should be considered carefully for applicability to other matters. Notably, this case involved a publicly available AI tool; courts may view enterprise AI platforms that maintain confidentiality and do not train on user inputs differently.

Related People



Meghan A. Podolny

Partner

+1 804 788 8458

mpodolny@hunton.com



Jessie Purtell

Associate

+1 804 344 7829

jpurtell@hunton.com

Related Services

Litigation

Artificial Intelligence

Technology

Discovery and E-Discovery

Media Contact

Lisa Franz

Director of Public Relations

Jeremy Heallen

Public Relations Senior Manager

mediarelations@Hunton.com

These materials have been prepared for informational purposes only and are not legal advice. This information is not intended to create an attorney-client or similar relationship. Case results do not guarantee or predict a similar result in any future case. Unless otherwise noted, attorneys not certified by the Texas Board of Legal Specialization. Hunton Andrews Kurth LLP is a Virginia limited liability partnership. Attorney Advertising.