



## Enforcement Insights

### SDNY Ruling Highlights Privilege Risks for Client-Generated AI Materials

April 15, 2026

In a February ruling in *United States v. Heppner*, Judge Jed S. Rakoff of the U.S. District Court for the Southern District of New York held that documents generated through a third-party generative AI tool and later shared with counsel were not protected by the attorney-client privilege or the work-product doctrine. The ruling underscores the potential discovery risks that can arise when individuals input confidential information into AI tools, including in the context of a potential or pending investigation or litigation.

The underlying criminal matter involves allegations of fraud arising from investments made by GWG Holdings, Inc. (GWG) into Beneficient to satisfy sham debts owed to a shell company, Highland Consolidated Limited Partnership (HCLP). Bradley Heppner—the founder of Beneficient, chairman of the GWG board, and controlling party of HCLP—was arrested in November 2025 on **charges** of securities fraud, wire fraud, conspiracy to commit securities and wire fraud, false statements to auditors, and falsification of records. In connection with their investigation, federal agents seized devices that contained thirty-one AI-generated documents that Heppner allegedly created using Anthropic's AI tool, Claude, to organize his thinking about the investigation, including potential defenses and legal arguments.

Heppner's counsel logged the documents as privileged, asserting that Heppner prepared them to synthesize his thoughts for communication with counsel. His counsel acknowledged, however, that Heppner created the documents on his own initiative, not at the direction of counsel. The government moved for a ruling that the AI-generated materials were not protected by the attorney-client privilege or the work-product doctrine.

**The district court agreed** with the government. First, the court held that the communications were not protected by the attorney-client privilege for three discrete reasons: (i) Claude is not an attorney, so there was no privileged communication between a client and counsel to generate the documents; (ii) Heppner's use of the consumer version of Claude undermined any claim of privilege because its policies put users on notice that the data could be disclosed to third parties without any confidentiality protections; and (iii) Heppner did not use Claude for the purpose of obtaining legal advice because he did not do so at the suggestion or direction of counsel. Second, the court rejected work-product protection claims because the documents were neither prepared by counsel nor prepared at counsel's direction, and did not reflect counsel's strategy.

The ruling is a clear caution for organizations and individuals using publicly available AI tools in connection with disputes, investigations, or other sensitive matters. Moving forward, clients should consider involving counsel before using generative AI for sensitive matters so that counsel can evaluate the specific tool and use, the applicable terms and policies, and any available safeguards to ensure documents and information remains privileged. Where AI use is appropriate, clients should consider using platforms designed to maintain confidentiality and should use those tools in a counsel-directed manner, recognizing that resolutions of disputes regarding privilege and disclosure may not be easily determinable and will likely remain highly fact-specific.