Questions to AI Models May Be Discoverable – IP Law Watch

On February 17, 2026 in *U.S. v. Heppner*, 1:25-cr-503 (S.D.N.Y., Feb. 17, 2026), Judge Rakoff held that a defendant's written exchanges with a public generative AI platform were not protected by the attorney-client privilege or the work product doctrine. The Government had seized approximately thirty-one documents memorializing the defendant's interactions with the public platform. Defense counsel asserted privilege because the inputs included attorney-learned information, were created to facilitate consultations with counsel, and were later shared with counsel.

The court found that defendant's communications were not protected by attorney-client privilege. First, the court found the communications were not between client and attorney. Second, the court found the communications were not confidential. The platform's privacy policy states that the model collects user inputs and model outputs for training and potential disclosure, including to governmental authorities and in connection with disputes or litigation, defeating any reasonable expectation of confidentiality. The court also found that platform's express disclaimer that it cannot provide legal advice undercut any claim that the exchanges were for the purpose of obtaining legal advice.

The court also found that defendant's communications were not protected by the work product doctrine. It found the AI-generated outputs were not prepared by or at the direction of counsel, nor did they reflect counsel's strategies; the defendant acted on his own when he created the documents.

Framing the issue as one of first impression, the court emphasized that AI's novelty does not alter settled privilege and work-product principles. Because the defendant's use of a public AI platform failed those tests, the documents were not protected.

To preserve existing privilege, publicly available AI platforms should be treated as third parties. One should assume user inputs and outputs may be retained, reviewed, and disclosed per provider policies and do not allow clients to share privileged facts or strategy with such tools. If AI must be used, route usage through documented, counsel-directed instructions, and only use tools offering enterprise-grade confidentiality terms. Litigation hold letters should ask whether any AI-initiated conversations relevant to the case exist to ensure practitioners get ahead of these issues. Lastly, communications with AI models before engagement or without counsel's direction should be carefully reviewed.

For additional analysis on this issue, please see this recent Litigation Minute.

By Chris Centurelli, Erik Halverson and Joshua Andrews