Gavel to Gavel: When "Ask AI" becomes Exhibit A

The hidden discovery risk of public AI tools

By: **Jake Krattiger** and **Tyler Self** // GableGotwals // March 11, 2026

Jake Krattiger

Tyler Self

Artificial intelligence is quickly becoming a default brainstorming tool, even in high-stakes legal situations. A recent federal decision makes clear that this convenience carries real litigation risk: what a user types into a *public* AI platform today may be discoverable tomorrow.

In *United States v. Heppner*, the U.S. District Court for the Southern District of New York held that a defendant's independent communications with a public AI platform were not protected by attorney-client privilege or the work product doctrine, even though the materials were later shared with counsel.

The ruling leaves little room for creative lawyering. Privilege does not arise simply because a litigant later loops in an attorney. Using a public AI tool is not the functional equivalent of consulting counsel, and passing along the output after the fact does not transform it into protected material.

The court applied orthodox privilege principles. Attorney-client privilege protects confidential communications between a client and a lawyer (or the lawyer's agent) for the purpose of obtaining legal advice. A public AI platform is neither. It owes no fiduciary duty and typically disclaims providing legal advice.

Confidentiality proved equally fatal to the privilege claim. The platform's terms allowed retention of user inputs, potential third-party disclosure, and the use of that data for model training. Those conditions undermined any reasonable expectation of confidentiality, a prerequisite to privilege.

The inherent purpose of the communication mattered. Because counsel did not direct the AI use, the defendant was effectively seeking advice from the platform itself, which falls outside the scope of privilege protection.

The work product doctrine offered no refuge. That protection applies to materials prepared by or at the direction of counsel in anticipation of litigation, particularly those reflecting counsel's mental impressions. The AI-generated analyses were created independently and did not reflect counsel's mental impressions when drafted. Later reliance by counsel did not convert them into protected work product.

Although decided in a criminal context, the reasoning translates to civil litigation, regulatory inquiries, and internal investigations. For companies, the message is direct: unsupervised use of public AI tools, especially after a subpoena, investigative demand, or credible threat of litigation, can create discoverable records that compromise privilege claims and litigation strategy.

AI governance is therefore a risk management function, not a tech preference. Clear internal policies, targeted training, and early legal oversight are not optional. They are the difference between leveraging AI effectively and handing opposing counsel an exhibit.

Jake Krattiger is a Shareholder at GableGotwals and a member of the Firm's Board of Directors, on which he serves as the vice president of talent development. Tyler Self is an Associate at GableGotwals, where his practice focuses on environmental and natural resources law.

Gavel to Gavel: When "Ask AI" becomes Exhibit A | The Journal Record