USA v. HEPPNER

AI, Attorney-Client Privilege, and the Risk to Law Firms

A Plain-Language Legal Analysis | February 2026

PART 1: KEY POINTS OF THE USA v. HEPPNER RULING

Background: What Happened?

Bradley Heppner was a financial executive charged with securities and wire fraud in a case involving an alleged $150 million scheme. After receiving a grand jury subpoena and hiring a defense team at Quinn Emanuel, Heppner used Anthropic's consumer AI chatbot Claude to help him think through his legal situation. He typed questions and information into Claude, generated 31 documents of back-and-forth prompts and answers, and then sent those documents to his lawyers.

Key Finding: Judge Rakoff ruled that the 31 AI-generated documents were NOT protected by attorney-client privilege or the work product doctrine — making them fully available to the government as evidence.

Key Point #1: An AI Is Not a Lawyer

For attorney-client privilege to apply, a communication must happen between a client and their lawyer. Judge Rakoff was clear: Claude is not a lawyer. It has no law license. It owes no duty of loyalty to the user. It cannot form a real attorney-client relationship. It is not bound by professional responsibility rules or confidentiality obligations.

Key Point #2: No Reasonable Expectation of Confidentiality

Anthropic's consumer privacy policy clearly states that the company collects data on user inputs and outputs, uses that data to train the AI model, and reserves the right to disclose data to third parties — including government regulatory authorities. When you voluntarily share information with a third party that does not promise to keep it secret, you lose your privilege over that information.

Key Point #3: Not Prepared at the Direction of Counsel

Heppner also tried to claim protection under the work product doctrine. This argument failed because his own lawyers admitted that Heppner created the documents 'of his own volition' and that his legal team 'did not direct' him to run the AI searches.

Key Point #4: Privilege Cannot Be Created Retroactively

Heppner argued that even if the documents were not privileged when created, sending them to his lawyers should have made them privileged. Judge Rakoff firmly rejected this idea. A document that is not privileged when created does not become privileged simply because it is later handed to a lawyer.

Key Point #5: A Possible Path Forward — the Kovel Doctrine

Judge Rakoff left a door slightly open. If Heppner's lawyers had actually directed him to use Claude as a tool — in the same way a lawyer might hire a financial expert or accountant — then the result might have been different. Under the Kovel doctrine, non-lawyer professionals hired by a lawyer can sometimes be included within the umbrella of attorney-client privilege.

PART 2: A WARNING FOR LAW FIRMS — WHEN STAFF USE AI WITHOUT PERMISSION

The Heppner ruling sent a clear message to clients: using consumer AI tools with sensitive legal information can destroy your legal protections. But what many law firms have not yet asked is this: what happens when the problem is not the client — but the firm's own staff?

What Is 'Shadow AI'?

'Shadow AI' refers to the use of AI tools inside an organization without the knowledge or approval of management. In a law firm, this practice can expose confidential client information and destroy legal privilege. According to technology researchers, the majority of employees in professional workplaces have used AI tools that were not officially approved by their employer.

Three Ways Unauthorized AI Use Can Hurt a Law Firm:

1. **Loss of Attorney-Client Privilege:** When a staff member types client information into a consumer AI tool without authorization, they may be sharing that information with a third party. Under Heppner, this could waive the privilege over not just the AI-generated content, but also the underlying client communications.

2. **Violation of Professional Responsibility Rules:** Lawyers have a duty of confidentiality under professional ethics rules. Rule 1.6 of the ABA Model Rules requires lawyers to make reasonable efforts to prevent unauthorized disclosure of client information.

3. **Malpractice and Client Claims:** If a client loses a case because a firm employee used an unauthorized AI tool, the client may sue the firm for legal malpractice.

Consumer AI vs. Enterprise AI: A Critical Difference

Consumer-grade AI platforms typically include terms of service that say the company may use your inputs to train the AI, store your data, or share it with others. Enterprise AI tools are different — they are sold specifically to businesses with contractual guarantees of confidentiality.

What Law Firms Need to Do Now:

1. Audit current AI use across the organization
2. Create a clear written AI policy
3. Train all staff — not just lawyers
4. Use only approved enterprise AI tools
5. Document AI use in client files
6. Update engagement letters and onboarding materials