Better Call Claude: The Confidentiality Risks of Getting Legal Advice from AI

Thomas McKay, Partner

Artificial intelligence is on the rise. With each passing week, it seems, there is another great leap forward in the capabilities of the newest AI tools. Many chatbots can now provide helpful answers to complex legal questions in a matter of minutes. As a result, people are increasingly turning to AI for legal advice, either in lieu of or in addition to consulting with lawyers.

Advances in AI present a great opportunity for efficiency in the legal industry. But the use of AI comes with risks. While much attention already has been drawn to incidents where chatbots hallucinate non-existent legal precedents, another major concern is the privacy implications of asking AI for legal advice.

The attorney-client privilege allows people to obtain legal advice free from the fear that their communications will be disclosed or used against them by an adversary. But these critical protections only apply to a person's communications with a lawyer, not with AI chatbots—no matter how much they may sound like lawyers these days.

A recent decision arising from a criminal prosecution in New York highlights the confidentiality risks that arise when a client asks AI for legal advice, and suggests some steps individuals can take to avoid these risks while still harnessing the power of AI.

The Attorney-Client and Work Product Privileges

Say you have a problem and you're not sure if you have done something wrong. Or you're planning to take some action and want to make sure you do it lawfully. Either way, you want to understand the legal implications. But you're afraid that if you tell someone about it, the information will be used against you some day. This is why our legal system recognizes the attorney-client privilege. People with such questions can consult with an attorney, and those communications generally will remain private so long as they were made confidentially and for the purpose of obtaining legal advice.

Now say you're in litigation or anticipate that you will be. You've hired attorneys. And you want to prepare some notes or documents that will help them help you. These may be privileged too. The work product privilege applies not just to an attorney's work, but also can protect documents that a client prepares for his attorney for the purpose of assisting in pending or anticipated litigation.

Together, these privileges play a valuable role in our legal system. They ensure that you can tell your lawyer the truth and obtain his frank advice. And they enable you to work collaboratively with your lawyer on your case, without fear that your thoughts or strategy will be disclosed to an adversary.

When AI tools enter the picture, however, these privileges are at risk.

A Cautionary Tale: *U.S. v. Heppner*

Bradley Heppner was chairman of GWG Holdings, Inc., a publicly traded company that filed for bankruptcy. Federal prosecutors began investigating Heppner for an alleged fraud and served him with a grand jury subpoena. Heppner hired lawyers from a prominent firm and began consulting with them about his case.

On his own initiative, Heppner also used Claude, an AI chatbot from Anthropic, to prepare several documents that contained legal and factual arguments relevant to his defense. Heppner then shared those documents with his attorneys.

The FBI later arrested Heppner and seized his electronic devices, which contained copies of the documents prepared with Claude. Federal prosecutors then filed a motion in court, arguing that the documents were not privileged and could be used against Heppner at his upcoming trial.

Heppner's lawyers argued that the documents were privileged. Although the lawyers had not instructed Heppner to prepare the documents, his purpose in doing so was to assist his lawyers. Plus, Heppner's queries likely reflected information he had learned during privileged conversations with his lawyers. A review of Heppner's queries and the ensuing responses very likely would give prosecutors insight into Heppner's privileged communications and trial strategy.

Nevertheless, in a brief oral ruling, Judge Jed S. Rakoff ruled that the documents were not protected by the attorney-client or work product privileges. The ruling relied on Judge Rakoff's view that Anthropic's then-applicable terms of service made clear that users have no expectation of privacy in their inputs: Anthropic generally collects data on user prompts, uses that data to train its AI tool, and may disclose that information to third parties including regulators. Given that, even if some of what Heppner entered in Claude was information he learned during privileged conversations, Heppner was waiving that privilege by effectively disclosing it to a third party. Similarly, even if Heppner was using Claude to organize his thoughts to help his attorneys formulate strategy, he did so by using a tool in which he could not have had a legitimate expectation of privacy.

As a result, federal prosecutors now can use Heppner's queries to Claude against him at his criminal trial, and may have gotten a preview of his and his lawyers' trial strategy.

Some Precautionary Measures to Take

Individuals and companies can avoid the same result by taking several precautions:

- *Be careful what you say to AI.* It can be tempting to think that conversations with a chatbot are private. As Heppner illustrates, however, they typically are not. So take care what you say. It's not the same as speaking to a lawyer. And even if your queries are entirely innocuous, adversaries in litigation often will find ways to use them to their advantage.

- *Know the terms of service and privacy settings.* Anthropic's privacy policy defeated Heppner's claim of privilege. Many widely available chatbots have similar policies. Although most chatbots now allow users to change certain privacy settings, it is best to assume that conversations with most consumer AI chatbots are not confidential. However, there are certain AI tools, such as the legal industry specific ones discussed below, which have privacy policies designed to protect privileged information.

- *Companies, train your employees.* As AI increasingly is used for everyday business functions, corporate employees may not realize that what they ask AI could someday be discoverable in litigation. Companies should train their employees about the confidentiality risks and develop policies about the appropriate uses of AI in the workplace. Those policies should include consideration of whether consumer chatbots may be used or whether to mandate use of industry-specific AI tools.

- *Ask your lawyer to use AI*. Many people ask AI for legal advice because it is quicker and less expensive than asking a lawyer, or because they've gotten a lawyer's advice but want a second opinion. These are understandable actions, but they present the risks discussed in this article. Many law firms are already using AI tools specifically designed for the legal industry. These tools not only conduct research and perform other legal tasks with great efficiency and at less cost to the client, but they take numerous steps to ensure that their use meets the privacy and confidentiality needs of the legal profession.

By asking your lawyer to use AI, individuals and companies can combine the speed and efficiency of a chatbot, with the confidentiality—not to mention the experience and judgment—of a lawyer.