Barclay Damon Live Presents *Cyber Sip™*
Season 5, Episode 3: "No Privilege for You, Heppner!—What Every Lawyer (and Client) Should Know Before Using AI"

Host: Kevin Szczepanski, Barclay Damon

[Kevin Szczepanski]: Are client's communications with an online AI platform such as Claude covered by attorney-client privilege or the work product doctrine? The answer, according to one federal district judge in the Southern District of New York, is "no." We're going to talk about *United States v. Heppner* today. And I think it's critically important that we do, because on one hand, generative AI is already becoming an invaluable tool, and lawyers and clients alike will insist that it become an indispensable tool in the years to come. And it may be very few years within which that indispensability becomes so critical to our practices. But along the way, we have to avoid some pitfalls that can get you and me and our clients in some fairly deep trouble when it comes to AI and this case, *US v. Heppner* is a fantastic lesson for us. Not so good what happened to Mr. Heppner and his counsel in the case, although it really was Mr. Heppner fault.

[Kevin]: But we're going to talk about all of that now. And the key takeaway is: clients do not input anything into an open online AI platform. Do not communicate with anyone who is not your lawyer, and make sure that if you are going to use any form of AI that you consult with and get the blessing of your lawyer before you do. All right, let's dive in. Heppner. What happens? Well, at a very high level, this is what the government alleges. Mr. Heppner is the executive of several corporations, including a holding company, and he is charged with defrauding the holding company's investors out of more than $150 million by making false representations about and causing the holding company to enter into self-serving transactions with at least two other businesses that he, Mr. Heppner, controls.

[Kevin]: So that's a problem. When does it become a problem? Less than a year ago. Let's go back to mid-2025. Mr. Heppner receives a grand jury subpoena. And the subpoena covers some of these very issues that we just talked about. He discusses it with his lawyer, but on his own, as the court in the Heppner case says, on his own volition, he goes on to a Claude AI platform. It's an open platform, meaning whatever you input trains the model and it's not confidential. And what does he do? Well, as you might expect, he's very concerned about the subpoena. And he does research that helps him—or at least he thinks it helps him—outline a defense strategy. He outlines what might be argued as to the facts and the law and what he thinks the government might be charging against him.

[Kevin]: So in his mind, Mr. Heppner is preparing these documents, these communications in anticipation of litigation. But he's doing it on an open AI platform. He's doing it without the knowledge of his lawyer and without the instruction of his lawyer. He's going it alone. And you and I both know what happens when you do that. All right. So what happens next? A few months later, it's late October 2025. A grand jury indicts Mr. Heppner on several charges... securities fraud, wire fraud, conspiracy to commit securities fraud and wire fraud, and other counts of fraud, including making false statements to auditors and falsifying corporate records. Within about a week in early November, Mr. Heppner is arrested and in conjunction with that arrest, the FBI executes a search warrant at Mr. Heppner's home.

[Kevin]: Here's where it gets interesting. FBI seizes numerous documents and electronic devices, and among

the seized materials are 31 documents reflecting communications between Mr. Heppner and the Claude AI platform. What happens next? Well, the government has these documents. Government shares them with Mr. Heppner's defense counsel, and Mr. Heppner's defense counsel does what any smart lawyer would do. He asserts the attorney-client privilege over all of these documents. But the government doesn't agree. The government asks the court... moves Judge Rakoff for ruling that these AI documents are not protected by either the attorney-client privilege or the work product doctrine. So Judge Rakoff has two issues to address.

[Kevin]: First, are these 31 communications with the Claude AI platform protected by the attorney-client privilege? And second, if they're not, are they nonetheless protected by the attorney work product doctrine? Judge Rakoff says no to both. Here's how he gets there. First, he discusses the attorney-client privilege. So we will as well. Judge Rakoff says the attorney-client privilege "attaches to and protects from disclosure communications between a client and his or her attorney that are intended to be, and in fact, were kept confidential, and that are made for the purpose of obtaining or providing legal advice."

[Kevin]: Let's take these three elements in turn. First, Judge Rakoff says, no, these communications Mr. Heppner had with Claude are not communications with his attorney. Why? Claude is not an attorney. Claude is an online software platform that provides AI services. It's not a licensed professional, so flat out cannot be considered an attorney. Second issue: were these communications intended to be and in fact kept confidential? Here again, Judge Rakoff says no, there's no reasonable expectation of confidentiality. And why is that? I'd like to read from this decision because I think it's really important, because on one hand you might think, well, Mr. Heppner is not a lawyer. He assumes that when he goes on Claude, the prompts that he's typing; the information he's putting into the model are not going to be shared with anyone. And in fact, maybe they never will be. But here's what the judge says. There's no confidentiality. And this is not merely because Heppner communicated with a third-party AI platform, but also—and here's the critical fact—"because the written privacy policy to which users of Claude consent provides that Anthropic" (that's Claude's parent company), "collects data on both users inputs and Claude's outputs."

[Kevin]: In other words, Claude uses the data that Mr. Heppner inputs to train Claude. And Anthropic reserves—I mean expressly reserves in its policy—that everyone including Mr. Heppner could see. It reserves the right to disclose that data to a host of third parties, including governmental regulatory authorities. Judge goes on: "the policy clearly puts Claude's users on notice that Anthropic, even without a subpoena compelling it to do so, may disclose personal data to third parties in connection with claims, disputes or litigation."

[Kevin]: In other words. Heppner could have had, quoting now "no reasonable expectation of confidentiality in his communications with Claude. The platform's very policy provides that Anthropic collects the data, and that it reserves the right to disclose that data to third parties." So not only is Claude not an attorney, but even if it could somehow be considered an attorney, the information that Mr. Heppner input into the AI platform is not confidential.

[Kevin]: So no attorney-client relationship, no confidential information, yet the judge goes on to address the third prong, which is: is Mr. Heppner, was Mr. Heppner seeking legal advice when he used the model? And the judge essentially says two things. First of all, as I said before, Claude is not a licensed professional. So Mr. Heppner could not have believed reasonably that he was seeking the advice of legal counsel. Second, judge notes that as part of the motion, the government itself submitted its own prompts to Claude, in which the platform said, I'm not a lawyer, I can't provide formal legal advice, and I recommend that you consult a qualified attorney who can assist you.

[Kevin]: So that too undermines the notion that, Mr. Heppner could have believed that he was obtaining or providing legal advice. And third, the court notes that Mr. Heppner did not make any of these communications

at the behest of his lawyer. Had counsel directed Heppner to use Claude, the judge says, "Claude might arguably be said to have functioned in a matter akin to a highly trained professional who may act as a lawyer's agent within the protection of the attorney-client privilege." But what matters for the attorney-client privilege, the court says, is whether Mr. Heppner intended to obtain legal advice from Claude. And here the record shows only that Mr. Hefner intended to—and eventually did—share with his attorney his communications with Claude. The court says, well, that doesn't matter, because the critical timeframe is the time in which Mr. Heppner is communicating with Claude.

[Kevin]: Is he obtaining legal advice at that time? If not, then it doesn't matter that he shares the communications with his lawyer later. Why does the court cut its so thinly? Because and you read this decision, it actually makes very good sense. The attorney-client privilege many people don't appreciate this is actually interpreted relatively narrowly. And the reason for that is that the privilege is intended, on one hand, to protect an individual and encourage the individual to communicate with his or her lawyer for the purpose of obtaining legal advice. But on the other hand, we'd never want the attorney-client privilege to shield facts or to prevent the government's right to conduct a criminal inquiry. So to that extent, the privilege is fairly narrow and the corollary to that is you can't shield public communications that you have with an individual or in this case, with an AI platform, by later sharing them with a lawyer.

[Kevin]: Communications do not become privileged just because they eventually get to your lawyer. So there's the court. One, two, three. Mr. Heppner doesn't communicate with a lawyer. His communications aren't confidential, and he didn't communicate with Claude for the purpose of seeking legal advice. Would it have been different if Mr. Heppner's lawyer had instructed him to communicate with Claude? Well, maybe. But that's just one of the three prongs. You would still have to deal with the fact that Claude is not an attorney. And secondly, that the platform Mr. Hefner was using was an open platform, and that Claude reserved the right to share the information that Mr. Heppner was providing with third parties.

[Kevin]: All right. So that's issue one. The court finds that Mr. Heppner's 31 communications with Claude are not privileged. But that doesn't end the day for Mr. Heppner, because if those communications are nonetheless work product of Mr. Heppner's attorney, they might still be protected. And the government might not get access to them. But as you know from reading the decision and as you'll hear next, Judge Rakoff says, no to the work product protection as well. Here's how he gets here. Judge Rakoff starts with the principle in place. The doctrine, the work product doctrine, provides qualified protection for materials prepared by or at the behest of counsel, in anticipation of litigation or for trial. Here's the problem for Mr. Heppner. The communications were not created by counsel. They were created by Mr. Heppner himself, and they were not created at the behest of counsel either, because at the time, Mr. Heppner was using Claude to make his queries and understand more about his potential defenses to the criminal charges that might come, he did not tell his lawyer that he was doing that. So because the communications could not have been created by or at the behest of counsel and the very purpose of the work product doctrine is to preserve the sanctity of counsel's mental thoughts, impressions, strategy, and so forth, the court says no, neither the attorney client privilege nor the work product doctrine protect Mr. Heppner's communications with Claude.

[Kevin]: So what can we take away from what we've learned from the Heppner case? I think there are a few takeaways, and they depend on who you are. If you're an attorney, you need to advise your client and make sure that your client understands that he or she should not be inputting sensitive information into an open AI platform, period. And he probably shouldn't be doing any writing or communicating with a third party without consulting with you first. So the advice should be whatever you do, do not communicate about your case with anyone, whether it's a live person or an AI platform. At the very least, check with me first. If you're an individual, same advice. Don't do it. I know you want to help your case. I know you're concerned. You're anxious, you're petrified about looming criminal charges. But if you go online and use Claude or ChatGPT or some other open AI platform, anything that you put into that platform, you are putting out to the world.

[Kevin]: So says Judge Rakoff. And why? Because we know that those platforms are not confidential. Don't

do it. And third, if you're a business, you want to have a carefully drafted AI policy so that your employees understand that when it comes to an open model like ChatGPT, Gemini, Claude, anything they input into the model is out for the world and they waive any potential protections, whether it's the attorney-client privilege, whether it is the work product doctrine, or even contractual protections for sensitive information. You should assume that anything you put into a public model is going to be discoverable. If it's in a civil case by your adversary, if it's a criminal case by the government and act accordingly. AI devices can be invaluable. And as we said at the top of this episode, I think they're going to become indispensable if we're not already there.

[Kevin]: And I don't think we're there yet. But when we're talking about the value and the power that AI can provide us, we're talking about enterprise or closed platforms. Anything that you are putting into an open platform, as Mr. Heppner did in his case, is out there for the world and is no longer protected by the privilege, by work product doctrine or any other legal protection you might otherwise have had.

[Kevin]: So word to the wise, be careful. Heed the lesson of the Heppner case, but by all means, continue to use AI safely to improve your efficiency, your effectiveness, and the affordability of your legal work. Because if you do it the right way, it only gets better. Hit me up in the comments if you thought this episode was informative. If you have other issues you think are important, if you think I got something wrong, anything, we really welcome your advice. Good to know you're out there, and it's good to know what you think about what we're putting out there. So thanks very much for joining us on this episode. We're back soon with another episode of *Cyber Sip*.

[Kevin]: The *Cyber Sip* podcast is available on barclaydamon.com, YouTube, LinkedIn, Apple Podcasts, and Spotify. Like, follow, share, and continue to listen.