Barclay Damon Live Presents *Cyber Sip*™

Season 5, Episode 4: "Recent Cases, Key Lessons: *Heppner* and *Warner* on Keeping Sensitive Data Out of Open AI"

Host: Kevin Szczepanski, Barclay Damon

[Kevin Szczepanski]: Hey, everyone, we are back with season five, episode four of *Cyber Sip*. And on our last episode, we talked about what terrible things can happen to you if you input sensitive information, attorney-client privileged information, sensitive documents, information about your legal case, your strategies, your strengths, your weaknesses into an open AI platform. And since our last episode, which we discussed *United States v. Heppner,* a second decision has come out. This time, the case of *Warner v. Gilbarco.* The *Heppner* case was in New York. The *Warner* case was in Michigan. And many people have been saying, look, we've got a new case *Warner*. It's very different. The courts are disagreeing, and so I wanted to have a second episode to address what bad things can happen when you put sensitive information into an open AI platform. And I wanted to talk to you about *Warner* because my secondary message is that *Heppner* and *Warner* are really very consistent. There are two ends of the same coin. One is a criminal case. One is a civil case. I'm going to talk about the differences between the two, and the key takeaways. But I'm not going to keep you in suspense. I'm going to give you the takeaway right now. If you are thinking about putting sensitive information, whether it's a communication with your lawyer, whether it's a prompt to identify the strengths and weaknesses of your case, whether it's something that you think might be bad for your case, and you're thinking of putting this into an open AI platform like Claude which can be open, and ChatGPT, which is nearly always open (although there are exceptions). Don't do it. At a bare minimum, talk to your lawyer first. He or she's going to tell you not to do it because if you do, you are taking a serious risk that your case can be harmed.

[Kevin]: In the case of *US v. Heppner,* the criminal defendant's case was harmed because he was forced to turn over the sensitive information that he input into Claude. In the case of Warner versus Gilbarco, which we're going to talk about today, the plaintiff, Warner, was not harmed. But do you really want to be in a position where a judge has to decide whether sensitive, secret information about your case is given over to your opponent? Of course you don't. So that's the takeaway. You can stop watching now, but please don't because we are going to talk about Heppner and Warner and how those decisions actually fit together. So let's start by giving you a little review of the US versus Heppner case. And I'm going to highlight the differences as we go on. And hopefully this will be a great conversation. Heppner is a criminal case and it was, it's being presided over by District Judge Jed Rakoff in the Southern District of New York. So in Heppner, here's what happens. The defendant, Heppner, receives a grand jury subpoena. And when he gets that subpoena and he talks to his lawyer about it, he realizes that he is the target of a federal criminal investigation. So what does he do? He does what you and I would do if we were in his position. He gets on the computer, he logs into Claude and he starts questioning Claude with prompts, like about the strengths and weaknesses of his case, what the government's theory of criminal liability against him might be, what his potential defenses might be. And here's the important part. His lawyer doesn't tell him to do this. And in fact, his lawyer doesn't even know about it. So what happens in that case?

[Kevin]: Judge Rakoff rules from the bench on February 10, just about 25 days ago. He rules that first, the prompts that Mr. Heppner has with Claude are not protected by the attorney-client privilege. Why? Because attorney-client communications need to be between a client and his attorney. And whatever else you might say about Claude AI, it's not a lawyer. That's number one. Number two, Judge Rakoff says these are not confidential communications. Why? Because they're communications with an open AI platform. And here's where I think it's important to read from the opinion. So it's not just that Mr. Heppner's communications are with a third-party AI platform, but as Judge Rakoff did, if you drill down into the written privacy policy and terms of use in Claude, what you find is number one, Claude's parent, Anthropic, collects data on both user inputs and outputs. In other words, Anthropic collects your prompts and it collects the work product that you get in response to your prompts.

[Kevin]: Number two, it uses those inputs and outputs to train Claude to become better. So, the information that you put in that you think is secret and sensitive is being used to train the model for everyone else in the world who uses Claude. And third, maybe most important or most nefarious, Anthropic reserves the right to disclose those prompts, inputs and outputs, to a host of what they call "third parties," including governmental regulatory authorities. So on that basis, Judge Rakoff says, well look, these are not confidential communications. In fact, Mr. Heppner could have had no reasonable expectation of confidentiality in his communications with Claude because according to Anthropic's privacy policy, in its terms of use, they're using that information to train the model. They reserve the right to give it over to third parties, including the government. Nothing confidential about that. Right? We'll talk about that a little bit more as we go on. And third, Judge Rakoff says, you know, even if these were communications with a lawyer and they were confidential, they weren't communications for the purposes of obtaining legal advice. That may seem counterintuitive. We talked about that in our last episode, but that's what Judge Rakoff said. He also went on to say that the communications that Mr. Heppner had with Claude were not protected by the work product doctrine either. Why not? Because the work...because work product has to reflect an attorney's mental impressions about the case and they have to be communications that were made by or at the behest of a lawyer. And in the case of Heppner, Mr. Heppner did not make the prompts to Claude at the behest of his lawyer, because his lawyer did not even know about the prompts, much less instruct Mr. Heppner to make them. And of course, as we know, Mr. Heppner himself is not a lawyer. In the Heppner case, Judge Rakoff says, nope, no attorney-client privilege, no work product. And so the secret communications that Mr. Heppner made to Claude that were discovered during a search by the FBI were not protected. The government gets them. Now, will they ultimately be admissible in the criminal case against Mr. Heppner? Maybe, maybe not. The judge can still find that even though the government gets to see them, they don't necessarily get to use them against Mr. Heppner in a criminal case. But unfortunately for Mr. Heppner, he's not going to find out about that until he's on the cusp of trial. So he is going to have those secret communications hanging over his head until his case is resolved. So that's US versus Heppner. Lesson? If you're a criminal defendant, don't put sensitive stuff into an open AI platform.

[Kevin]: But interestingly enough, on the very same day that Judge Rakoff ruled, February 10, 2026, there was another decision coming out in the Eastern District of Michigan. And that decision came in the case of Warner versus Gilbarco, Inc. And unlike Heppner, Warner is a civil case. And so that's going to be important, as we'll find out, because as a civil case, it is subject to the federal rules of civil procedure on discovery. Heppner, documents come to light during the execution of a search warrant. Warner, the documents come to light during civil litigation. So the civil rules apply. Here's what happened in Warner. Let me give you a rundown of the facts. Lots of discovery motions. The plaintiff is *pro se*, which for you non-lawyers who are not familiar with all of our legal jargon, means Ms. Warner is representing herself. Now you might say, well, she's kind of like Mr. Heppner then. She's just a plaintiff, just as he was a defendant. Not so, because Ms. Warner is *pro se*, meaning that she is acting as her own lawyer. So for purposes of the legal analysis, she's a lawyer. She's a lawyer and a client. What's that old adage? Person who represents themselves has a fool for a client. Well, in this case, you're going to assume here that Ms. Warner was no fool. She actually handled herself very well. So the court sets a discovery deadline. There is a flurry of discovery motions going back and forth. For my non-lawyer friends, that means lawyers are fighting with each other about whether they get this or that document, whether they get to ask this or that question under oath. And this is not the most hotly contested case I've ever seen, but it's a hotly contested one. So the defense lawyers and the defendants are represented by lawyers. They are, in the words of, Judge Patti, who's the magistrate judge presiding over the Warner case, they are, and I want to make sure I get this right, they are, in Judge Patti's words, "preoccupied." They're a bit preoccupied with the plaintiff's use of AI, and according to Judge Patti, that preoccupation needs to stop, or in the words of Judge Patti, it needs to "abate." Why use "stop" when you can use "abate"? Anyway. It's a contested case and Ms. Warner is using AI, this case, ChatGPT, a fair bit. And she is, why is she doing that? Because she's not a lawyer. So she's using the best thing, the best tool she can find to ask about strengths and weaknesses of her case, ask about issues that come up in the case, very similar to what Mr. Heppner was trying to do in his case.

[Kevin]: So, here's what happens. The discovery deadline passes. The court's got a deadline saying you must complete discovery by December 5 of last year, December 5, 2025. And according to the rules, you had to not only finish discovery, but you had to make a motion. That is, if you were going to ask the judge for anything else you needed that the other side hadn't given you, you needed to do that by December 5 as well. Well, the defendants break that rule, they don't bring their motion requesting more stuff from the plaintiff until several weeks after the discovery deadline. So the court's decision starts by noticing, you know, I sent out a discovery order. I asked you to complete discovery by this date. It's more than an ask, it's an order, and you didn't do it. And not only that, but I have court rules and the district judge in the case, Judge Drain, Gershwin A. Drain, to whom the magistrate judge reports, he's got his own rules saying you have to do this before the discovery deadline expires. So if you're reading Warner, you're thinking the magistrate judge who wrote the opinion was really kind of miffed at the fact that the defendants were making this motion late.

[Kevin]: And that might well have played into the judge's ruling. But to Judge Patti's credit, he did go on to address the merits of the case. So here's what happens. He says, look, it's late, but let me focus on this request. The defendants are, I don't want to say "obsessed." Let me use Judge Patti's words. They were "preoccupied" with the plaintiff's use of AI. So they made a motion. And they wanted all documents and information relating to the plaintiff's prompting of ChatGPT and the outputs or work product she received from ChatGPT. And so the judge had to wrestle with this issue. But unlike Judge Rakoff in the Heppner case, Judge Patti had to wrestle with it under the federal rules of civil procedure. And I want to read from the rule just a little bit so that we can have that context. So this is federal rule of civil procedure 26B and it's subsection 1. And it states that the parties "may get discovery regarding any non-privileged matter," that's not relevant here, "that is relevant to any party's claim or defense and proportional to the needs of the case." What does that mean? That means the court is going to weigh the party's interests to determine whether it's really fair and reasonable for one party to ask the other party to produce something. So what are the factors to be considered? Well, "the importance of the issues at stake in the case, the amount in controversy, the party's relative access to relevant information, their resources, the importance of discovery in resolving the issues, and finally, whether the burden or expense of the proposed discovery outweighs its likely benefit." Now there's a lot there, but I want to focus on one particular portion of that balancing test, and that is the party's relative access to relevant information. And that's going to become key here because as you will see, the judge is mindful of the fact that the defendants want the plaintiff's secret prompts and outputs from ChatGPT.

[Kevin]: But the defendants don't really need that in order to defend themselves in the case. Why is that? Because there are discovery tools the defendants can use to get the information they need. They have, they can make document requests to get the written documents they need. They can serve what are called interrogatories or written questions that the defendants have or the plaintiff has to answer in order to get more information. And one of the most valuable tools that the defendants have at the ready is the tool of the deposition. That's where they get to sit the plaintiff down, and other witnesses down, to ask them questions under oath. And you've seen that on television and court TV, the ability to cross-examine a witness, to probe the strengths of their claims, the strengths of their defenses is pivotal to the litigation of any case. Defendants can do all of that without having to have access to the plaintiff's AI prompts and the responses that she got from ChatGPT. So here's what the court says in light of this. Look, first of all, there's no evidence that the plaintiff uploaded any confidential documents. And the court started with that because there was a confidentiality order in place and the judge, Judge Patti, wanted to satisfy himself that the plaintiff hadn't done anything to violate the confidentiality of any document. And second, in all likelihood, wanted to satisfy himself that the plaintiff hadn't waived the confidentiality of any particular document. So Judge Patti starts by saying, look, there's no evidence of that at all.

[Kevin]: Second, these documents were clearly prepared in anticipation of litigation. In fact, they were... these AI prompts, and their outputs were created during the heat of the litigation. There's no question that they were prepared to aid, however they might have, aided the plaintiff in her prosecution of her claims. So then the judge says, look, these are not relevant. And even if they are, the discovery of these AI prompts and responses are not proportional under that federal rule I just read you. Now, let's break that down. Let's talk about the relevance because the defendants argued that these AI prompts and responses were relevant. And here's what the defendant said in their motion, their request to the court to compel the plaintiff to produce this stuff. Here's how they're relevant. They're directly relevant to the issues that include at a minimum... Already we can see the defendants are gilding the lily here. That's an old term, but I rather like it. Here's how they're relevant. The authenticity and authorship of filings and assertions made in the case, potential admissions against interest embedded in her prompts. For example, defendant even gives us examples. She could have identified weaknesses in her claims and position. She could have described those weaknesses and asked the third-party generative AI tool, otherwise known as ChatGPT, to address and minimize those weaknesses. Or she could have admitted the elements of her *prima facie* claims. Third, whether she disclosed or publicized information pertaining to the case to a third party, her preservation of evidence, that's number four, and five, her credibility. So that's the defendant's argument. I've got to tell you, it's pretty thin. It's thin because it's vague and it reeks of the kind of speculative argument that courts just hate in the context of discovery. It sounds like the defendants want to go on a fishing expedition. It's almost as though they're saying, hey, judge, we know that the plaintiff has researched her claims on AI and we just have a feeling that there's something really good there that we could use. So you need to let us have it. And the judge is saying, no, I don't see that any of this is relevant or if it is relevant, it's not proportional. And the judge doesn't say this in so many terms, but it makes sense. And why is it not proportional? It's not proportional for the very reasons we just discussed a moment ago. The defendants have access to written discovery, interrogatories—that's written questions, depositions—that's questions under oath. There are so many different discovery tools that the defendants could have used to gather the proof they needed to disprove the plaintiff's case and establish their own defenses. And the fact that they waited until after discovery was over to guess or speculate about what good stuff the plaintiff might have in her AI prompts and responses is not good enough. It just doesn't pass the relevance or proportionality thresholds that are set out in that federal rule of civil procedure 26B.

[Kevin]: The court takes a step further. The court says, okay, I noticed, defendants, that you also asked me to order Ms. Warner to produce a privilege log so that she can put, she can disclose in effect the dates, the subject matter of all of her AI prompts and responses. And then we can decide whether she should have to turn them over. And the judge said no. And the reason the judge said no is because he concluded that the work product doctrine applied. Why did the work product doctrine apply? Well, first of all, the work product doctrine has to involve an attorney's mental impressions. And the court didn't say this in so many words, but clearly a *pro se* plaintiff who's acting as her own lawyer, prompting ChatGPT for information about strengths, weaknesses of her claims, strengths or weaknesses of the defendant's defenses to those claims relates to her mental impressions. And secondly, the communications were made by or at the behest of a lawyer because Ms. Warner is acting as her own lawyer. The court also addressed, last point, the issue of waiver. Okay, so this is work product because it's all about Ms. Warner's mental impressions of her case and Ms. Warner is acting as her own lawyer, so it's an attorney's mental impressions in effect, but she put all this stuff into ChatGPT. So whatever protection she had under the work product doctrine was waived, poof, it's gone. And the court addressed that as well and in a very interesting way. Judge Patti said, you know, the work product waiver, in order to waive that work product protection, it has to be a waiver to an adversary or it has to be a waiver that is likely to get into an adversary's hand. And here, Judge Patti uses some shorthand, but I think we can glean what he means from it. He goes on to say, and ChatGPT and other generative AI programs, are tools not persons, even if they may have administrators somewhere in the background. So in other words, it's not a waiver because Ms. Warner did not disclose her work product to an adversary. She disclosed it to an AI platform and reading between the lines, she didn't do so in a way that was likely to get into an adversary's hand. Two points to that. Number one, when you're uploading a prompt or you're prompting an open AI platform, there's no direct line from your prompt into the hands of some other person or individual, least of all your opponent in a civil case. Second, think of how difficult it would be for a defendant to use ChatGPT to find information about your specific prompt and the specific outputs that those prompts created. Not so easy, right? I mean, I'd have to be thinking pretty carefully. I'd have to almost find a needle in a haystack to figure out what it is you or Ms. Warner uploaded or input into ChatGPT. Now, Judge Patti doesn't say that, but clearly that's the import of it. In that sense, it's not a disclosure to an adversary. It's not...in Judge Patti's words, likely to get into an adversary's hands. And how do we know that? We know that by the very nature of the defendant's motion. The defendant goes to court and makes a motion asking the court to force the plaintiff to turn over her AI prompts and her AI outputs. If that information had already fallen into the defendant's hands, they would have no reason to ask the court to force the plaintiff to turn it over. So in that sense, the court appears to have the comfort that merely prompting ChatGPT and getting outputs from ChatGPT does not constitute a waiver of the work product doctrine. So in the Warner case, the judge denies the defendant's motion and refuses to direct the plaintiff to turn over her AI prompts and the outputs that resulted.

[Kevin]: All right. With all that said, let's talk about the key differences between the Heppner case on one hand and the Warner case and the other, why these cases are different enough that they can be read together. They're not inconsistent. We're going to harmonize them now. First off, civil versus criminal. In the Heppner case, you've got a criminal case against a criminal defendant and the AI prompts and results that are turned over are turned over in the course of a search warrant. The Warner case on the other hand is a civil case, and the court is determining whether the defendants are entitled to the plaintiff's AI information in the context of federal rule of civil procedure 26B. And that rule, as we've discussed, requires a balancing test. Judge Rakoff is not required to balance the party's interests in the Heppner case, but in the Warner case, Judge Patti is. That's one big difference. Next big difference, and I'm not sure how big it is, but it is a difference and we'll never know whether it played into the decision, but I think it did.

[Kevin]: In the Heppner case, there's no discovery deadline. There's no issue about the timing of request of Judge Rakoff that he hold the defendant's AI information inadmissible. In the Warner case, however, there was a discovery deadline. The defendants missed it and they didn't ask the court for the plaintiff's AI materials until that deadline passed. Now, as I said earlier, Judge Patti did go ahead and rule on the merits of the defendant's request. So we'll never know whether missing that deadline played into Judge Patti's decision. But if you read the decision, it does appear that the lateness of the defendant's request was a factor. Well, here's a big factor. Third factor is the work product doctrine. Now in the Heppner case, it was the defendant himself that prompted Claude AI for legal information. The defendant's lawyer did not tell Mr. Heppner to do it. In fact, defendant's lawyer didn't even know that Mr. Heppner was doing it. So it was relatively easy, I think, for Judge Rakoff to conclude that even if these communications involve the mental impressions of Mr. Heppner's lawyer, let's assume that they were not made by a lawyer or at the behest of a lawyer. So the work product doctrine does not come into play. But in Warner, the analysis is very different. Even though it was the plaintiff herself who was prompting ChatGPT, the plaintiff was *pro se*. She was representing herself and therefore acting as a lawyer. And so her communications with ChatGPT were made by or at the behest of a lawyer because she's acting as her own lawyer. So based on those three factors, I think you can see that Heppner and Warner involve very different contexts, very different facts. And so they're not inconsistent with each other. They appear to be the right rulings based on the situation that was before Judge Rakoff and Heppner, Judge Patti and Warner. And so we can harmonize them. We can read them together. But let me close with this.

[Kevin]: I do think that however we can harmonize Heppner and Warner, there is an issue that will flow out of both of these decisions, and that is the question of confidentiality. When Judge Rakoff found that the communications Mr. Heppner had with Claude AI were not confidential, he specifically noted that Mr. Heppner was in effect communicating with a third party, Claude AI, and the privacy policies and terms and conditions of Claude AI, whether he knew them or not, made very clear that the information wasn't confidential, it was going to be used to train the model, and it could be disclosed to other third parties, including the government.

So in those circumstances, if you read Heppner, you might think that Judge Rakoff or more broadly, the courts in New York would find that a communication between an individual or even a lawyer with an open AI source is *per se* not confidential. In Warner, it reads a little bit differently. The court goes out of its way in talking about work product to say, know, yes, Ms. Warner was communicating with ChatGPT, but that's not an adversary. That is just a platform. And what's more communicating with the platform has to make it likely that an adversary will be able to get his or her, its hands on that sensitive information. And is that really likely? And Judge Rakoff seems to think so. And he reads the privacy policy, the terms of use of Claude AI and says, hey, you know, if they're saying that they're training the model and they're saying that they can turn this over to third parties, how confidential is it? Judge Patti doesn't go this far, but you can see the argument coming. How likely is it that my prompting an open platform is going to result in the substance of that prompt finding its way into the hands of my adversary in some civil litigation. How would my adversary even know to look for that information? My adversary would have to guess that I'm using AI. We have to further guess what platform I'm using, assume that it's an open platform, and then be nimble enough to prompt that platform for information that I myself put in there. I'm not sure how likely that is. And it gets even more difficult when you think about your adversary being able to prompt the platform to obtain outputs that you might have received.

[Kevin]: I think Heppner and Warner are very interesting cases. We need to pay very close attention to them. And I am predicting that down the road in the next 12 to 18 months, we're going to see cases on both sides and the question that will be litigated is whether a party really does lose confidentiality by inputting data into an open AI platform. But let's end with this. Do you really want to leave it in the hands of a court to tell you whether you've lost the attorney-client privilege, whether you've lost the protection of the work product doctrine just because you decided to do research on your own as a client or as a lawyer and in doing so, you put sensitive information into Claude or ChatGPT or Gemini? You don't. So, Heppner and Warner are helpful to us lawyers in setting the guardrails or the contours of this legal issue. And if something bad happens, we know what we can argue after the fact to try to preserve the protection of the privilege or the work product doctrine. But in the end, you don't want to be there. And so as I started, I think the best advice for now is to reserve that sensitive research for closed platforms, platforms like Lexis[Nexis] Protégé. There are other closed platforms on ChatGPT itself. And if you review the terms of use and the ChatGPT website, you can identify them. There are still other, many different types of closed platforms that you can use to conduct research that is not used to train the model and is not going to be accessible to the public.

[Kevin]: But aside from that, don't use ChatGPT, don't use Claude. And if you're ever tempted to do so, talk to your lawyer first because he or she will give you some really helpful advice that'll either talk you out of it or help you protect yourself before you waive the privilege or work product or the other protections that you have as a citizen of the United States. So I hope this is helpful. Please hit me up in the comments with any questions or thoughts you have; like, comment, and share. I've enjoyed bringing this to you. These are fascinating cases. And I think even though they're kind of specific in-the-weedsy, lawyer-y kind of issues, they're really issues that apply to all of us, because in the end, we're just talking about locking the door, locking the safety deposit box and not leaving valuable things out in the open. So I hope this was helpful and we'll be 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.

This material is for informational purposes only and does not constitute legal advice or legal opinion. No attorney-client relationship has been established or implied.

Barclay Damon Live *podcast transcripts and captions are automatically generated through artificial intelligence, and the texts may not have been thoroughly reviewed. The authoritative record of* Barclay Damon Live *programming is the audio file. Thanks for listening.*