

Do You Waive Privilege by Using AI? Two Federal Courts Say It Depends

Introduction

Welcome to the Case of the Week segment of the Meet and Confer podcast. My name is Kelly Twigger, and this week we are going to talk about not one, but TWO crucial rulings that were just issued in February 2026 that are the first rulings on the question of whether materials generated using publicly available AI tools are discoverable, and under what circumstances those materials are protected by the attorney-client privilege or the attorney work product doctrine.

I'm recording today's episode just a few days after concluding the University of Florida eDiscovery Conference. If you were one of the more than 4500 people who joined us for the eDiscovery event there, you heard us discuss these rulings at length on both my case law panel with [Maria Salacuse](#) from the EEOC and [United States District Judge Xavier Rodriguez](#), as well as several other sessions at the conference. And as usual with discovery decisions, there is much to discuss.

One of Judge Rodriguez's points from the panel on these decisions was key – and that was that viewing these decisions through the lens of the law as it exists today, both rulings are correct, but the question is whether there needs to be a rethinking of the law on privilege in the age of AI.

If you've been following this issue on LinkedIn, you saw an absolute EXPLOSION of discussion over United States District Judge Rakoff's oral decision in the *U.S. v. Heppner* case that was issued on February 10th, and then his written decision that was issued on [February 17th](#). But what you didn't see much at all on was United States Magistrate Judge Patti's decision in *Warner v. Gilbarco*, issued on [February 10, 2026](#), the same day as District Judge Rakoff's oral ruling. And that's where we're going to start today. Several firms out there issued client updates on these decisions following our discussion at UF, which is fantastic. Everybody needs to know about this. Most of them start with *Heppner*, and that makes sense in our minds because Magistrate Judge

Patti's decision seems to, without stating it, really reject what District Judge Rakoff says. But Magistrate Judge Patti's written decision came first.

As always, we'll provide links to the decisions in Minerva26 for both cases, and you can use the Generative AI issue tag in the platform to be immediately notified of any new decisions on Generative AI if you're a subscriber.

These are both very short decisions, and I suggest you read them, distribute them, talk about them in your firms, in your in-house counsel departments, and what they mean for you. They are going to impact how you advise your attorneys and your clients on the use of AI. And as we're going to discuss, there are still a lot of questions out there.

Let's take a look at these cases.

Warner v. Gilbarco

Let's start with, as I mentioned, *Warner v. Gilbarco*. The parties were before the Court in the *Warner* case on multiple motions, and our discussion today focuses on the defendant's Motion to Compel and for Sanctions. Each of the motions at issue was referred to Magistrate Judge Patti for resolution by the District Court.

Defendant's motion sought to have the Court overrule plaintiffs' attorney-client privilege and attorney work product doctrine objections to AI materials that she created based on a theory that she waived the privilege by using a publicly available AI tool, or in the alternative, to require her to produce a privilege log listing those materials. The Court's decision is a bit hard to decipher. It's all one long paragraph in one bullet point under section two on the defendant's motion. So it's a bit to piece through, but Magistrate Judge Patti says that the motion seeks "intrusive post-discovery production based on speculation about what might exist in Plaintiff's internal drafting process."

Plaintiff here was a pro se litigant – meaning she was representing herself – who used ChatGPT, the AI platform from OpenAI, for her own internal analysis and mental impressions in assisting with her case. She contended that the interactions were her thought processes, not an existing document, and were not discoverable both on that basis and also that discovery had closed in the matter.

Magistrate Judge Patti denied both requests from Defendants, finding first that the AI materials created by the plaintiff were not discoverable, and that even if they were, they were protected by the attorney work product doctrine because she was permitted to

assert because the plaintiff was pro se.

The next part of Magistrate Judge Patti's ruling is key. In addressing the defendant's argument that plaintiff waived the privilege by disclosing information to ChatGPT, the Court found that waiver has to be "to an adversary or in a way likely to get in an adversary's hand." What the Court says next is the most critical piece of the analysis as to the role that AI tools will play in the litigation process:

And ChatGPT (and other generative AI programs) are *tools*, not persons, even if they may have administrators somewhere in the background. As plaintiff noted in her response, Defendants' motion "asks the Court to compel Plaintiff's internal analysis and mental impressions—i.e., her thought process— rather than any existing document or evidence, which is not discoverable as a matter of law. The motion seeks intrusive post-discovery production based on speculation about what might exist in Plaintiff's internal drafting process, untethered from Rule 26 relevance, disregarding the heightened protection afforded to opinion work product, and improperly attempting to manufacture a waiver where none exists. At its core, Defendants' request is a fishing expedition. ... (emphasis added)

A lot of that language, frankly, was pieced directly from the plaintiff's brief. The Court went on to say – and again, this is very key – that it agreed with plaintiff that defendant's "pursuit of this information is 'a distraction from the merits of this case[,]'" and that Defendants' theory, which is supported by no case law but only a Law360 article posing rhetorical questions, 'would nullify work-product protection in nearly every modern drafting environment, a result no court has endorsed.'"

And that is it from Magistrate Judge Patti in *Warner* on the issue. But what he leaves us with is two things – FIRST, that the content – what I assume includes both the prompts and the output, the generated content that plaintiff created – are NOT discoverable. There's little explanation as to why other than that they are post-discovery production, but he does seem to say that they are not a document, as we would require under Rule 26. Now remember that Rule 26 applies to electronically stored information, which can be a document and documents. It's a separate part of that query. So is AI generated content the output? Does it include the prompt? Are those ESI that are discoverable? That seems to have been decided in the *OpenAI* cases.

So it's a little bit of a question as to how Magistrate Judge Patti is really addressing that here. There is a separate inquiry, though, that this appears to be post-discovery – after the close of discovery – and that's what makes it not discoverable in this context. What

that means for us is that you have to pay attention to the facts. And just as we have always discussed here on Case of the Week, the facts—what's really crucial to the analysis—are what you have to focus on. And we're going to get to some more of that as we get through our discussion here.

Magistrate Judge Patti also leaves us with this question of how do we deal with talking to custodians now about their AI-related content? If we're still in the open discovery period, do we now have to have conversations with our custodians and our clients about what AI generated tools that they use or that their custodians may be using to generate content? And we have to have frank discussions with those custodians about how they're used. That takes us into the conversation about whether those prompts are preserved, where they're preserved, what is accessible, are these publicly available AI platforms accessible? Or are we talking about using Microsoft Copilot within Teams? So it's so many questions that these decisions raise.

Second, Magistrate Judge Patti leaves us with the statement that AI platforms like ChatGPT, and likely Claude, Gemini, and others are TOOLS, they are *not people* and the disclosure of information in litigation to the tool does not constitute waiver because a tool is NOT a person. He goes on further to say that the theory of waiver in this context would “nullify work-product protection in nearly every modern drafting environment, a result no court has endorsed.” And that is really critical here because even as I'm typing out information in a Google Doc or in a WordPerfect document, that system is autofilling my paragraph. Does that make it AI generated content? How are we defining what is AI generated content? What is discoverable? When are our communications with these tools going to be discoverable? These are questions we don't have the answers to yet.

What Magistrate Judge Patti did not discuss in the *Warner* decision was the confidentiality of the platform that plaintiff used in *Warner* or whether she was using a free or a paid version, as well as what the privacy policy is from OpenAI for ChatGPT for the version she was using. And that's really critical as we transition into talking about the Rakoff decision.

U.S. v. Heppner

That exact same day that Magistrate Judge Patti issued his written decision in *Warner*, United States District Judge Rakoff ruled from the bench in a criminal case titled *U.S. v. Heppner* also on the privilege issue, but on a VERY different set of facts. District Judge Rakoff's ruling held that AI generated content created by a criminal defendant and shared with counsel was not protected by the attorney-client privilege or the work

product doctrine. The defendant in *Heppner* was under indictment for fraud to the tune of more than \$150 million. In his oral ruling, District Judge Rakoff dismissed any claim of privilege on the grounds that the defendant could “disclose” it to a third-party, in effect AI, which had an express provision that what was submitted was “not confidential.” That oral ruling, to many of us, and I say many because the chatter on LinkedIn was absolutely rampant, meant that District Judge Rakoff was saying AI was a person. And that’s problematic, and completely contrary to what Magistrate Judge Patti found.

A week later, on February 17th, District Judge Rakoff issued a thoughtful written decision reaching the same conclusion, but not addressing the “is AI a person” argument. The defendant was indicted on multiple securities fraud charges on October 28, 2025. The indictment was unsealed on November 4, 2025, and the defendant was arrested a day later. As part of the arrest, the FBI seized multiple electronic devices from the defendant's home. Following the seizure, the defendant's counsel advised the government that there were 31 documents on those various devices that memorialized communications that the defendant had had with Claude – the Generative AI platform from Anthropic, which has also been in the news lately for different reasons. According to counsel, the defendant created the materials with Claude after it was clear he was the target of an investigation, but without any suggestion from counsel that he do so. The defendant used Claude to “prepare reports that outlined defense strategy, that outlined what he might argue with respect to the facts and the law that we anticipated hearing based on what the government might be charging.” Counsel argued that the reports were prepared in anticipation of his indictment and were privileged.

The defendant asserted that he had input information he received from counsel, that he created the documents for the purpose of talking with counsel to obtain legal advice, and that he shared the contents of the documents with counsel. But counsel conceded that they had not directed the defendant to run searches using Claude.

So in *Heppner*, we've got the defendant essentially conceding that these are documents, potentially because he took them out of the platform and created actual documents and some other electronic means, whether that be Word, Google Docs, etc. We don't know if they were still just a prompt with an output in an AI tool on his desktop. In process, the government segregated the documents in *Heppner* and then made a motion for a ruling that the documents were not protected by the attorney-client privilege or the work product doctrine.

As to the attorney-client privilege, the Court found that the discussions were with Claude, which is not an attorney, and that alone was enough to determine no privilege existed.

Next, the Court really specifically addressed the confidentiality of defendant's communications with Claude and found that there was no confidentiality—not just because he communicated with an AI platform, but because the written privacy policy that defendant consented to as a user provides that defendant's content may be used to train Claude and that Anthropic reserved the right to disclose such data to third parties including government regulatory authorities: “The policy clearly puts Claude's users on notice that Anthropic, even in the absence of a subpoena compelling it to do so, may ‘disclose personal data to third parties in connection with claims, disputes[,] or litigation.’”

And it ignores the Stored Communications Act which has been leveraged by every technology company to preclude providing user-generated content except where the government issues a warrant for it, which certainly could have happened here because we're in the criminal context, but not in the *Warner* decision which was civil. There are also multiple other questions here as to whether or not all of the outputs continue to be retained, whether they're even available. There was a lot of discussion about retention in the *OpenAI* cases and the ability for that content to be retained. So whether it's even going to be available for anyone to subpoena is an open issue.

As I mentioned, District Judge Rakoff relied on two of the dozens of rulings in the *OpenAI* case located [here](#) and [here](#) to find that “AI users do not have substantial privacy interests in their ‘conversations with [another publicly accessible AI platform] which users voluntarily disclosed’ to the platform and which the platform ‘retains in the normal course of its business.’”

District Judge Rakoff also found that the defendant did not communicate with Claude for the purpose of obtaining legal advice because he did not do so at the advice or suggestion of counsel. According to the Court:

Had counsel directed Heppner to use Claude, Claude might arguably be said to have functioned in a manner akin to a highly trained professional who may act as a lawyer's agent within the protection of the attorney-client privilege. But because Heppner communicated with Claude of his own volition, what matters for the attorney-client privilege is whether Heppner intended to obtain legal advice from Claude, not whether he later shared Claude's outputs with counsel. And Claude disclaims providing legal advice. Indeed, when the Government asked Claude whether it could give legal advice, it responded that “I'm not a lawyer and can't provide formal legal advice or recommendations” and went on to recommend that

a user “should consult with a qualified attorney who can properly assess your specific circumstances.”

District Judge Rakoff said “it is black-letter law that non-privileged communications are not somehow alchemically changed into privileged ones upon being shared with counsel.”

The Court similarly dismissed the applicability of the work product doctrine to defendant’s Generative AI materials because they were “not ‘prepared by or at the behest of counsel,’ nor did they reflect defense counsel’s strategy.” Ironically, they were positioned by defense counsel as absolutely having defense counsel strategy in them, but one that was concocted by the defendant himself. The Court found that the defendant was not acting as his counsel’s agent when he communicated with Claude and therefore the privilege did not apply.

In denying the application of either privilege to the AI generated materials, District Judge Rakoff granted the government’s motion. But he left us with this, that again begs the question that Judge Rodriguez raised of whether we need a new standard for dealing with AI generated content:

Generative artificial intelligence presents a new frontier in the ongoing dialogue between technology and the law. Time will tell whether, as in the case of other technological advances, generative artificial intelligence will fulfill its promise to revolutionize the way we process information. But AI’s novelty does not mean that its use is not subject to longstanding legal principles, such as those governing the attorney-client privilege and the work product doctrine. Because Heppner’s use of Claude fails to satisfy either of these rules, the AI Documents do not merit the protections Heppner has claimed.

These two cases present very different factual scenarios asking the same question and giving us different answers. There is the notion that the standard in criminal cases is often different than that of civil cases, but we are applying the same principles of attorney-client privilege and attorney work product doctrine to both. As of today, neither decision has been appealed, so we’ll keep an eye on both cases. Neither of the transcripts from the hearings was available yet, so those may help provide more context. In fact, Magistrate Judge Patti even noted in his decision that for the reasons he was stating in his written ruling, but also for the reasons that were provided at the hearing. So it would be telling to understand more of what he said in that hearing.

Takeaways

What do these decisions mean for litigators and how do we advise our clients? I think that most everyone agrees that's a very tricky question to answer right now. Magistrate Judge Patti's decision protected the AI materials, but the context of them was unclear. They were post-discovery, were they impressions? Were they documents? If so, they would not have been. Was discovery closed when the plaintiff generated them? If so, they would not have been discoverable. Was there any content that she generated before discovery was closed? Was there another reason?

Is AI a person as District Judge Rakoff suggested in his oral ruling or is it a tool as Magistrate Judge Patti stated outright in his written decision? District Judge Rakoff did not address that again in his written ruling. My personal opinion is that courts will come down on the view that AI platforms are a tool, and that Magistrate Judge Patti's language that to consider otherwise "would nullify work-product protection in nearly every modern drafting environment" reflects practical reality. After all, in the world of software development, if you aren't integrating AI, you are at risk of becoming obsolete very quickly. This is a technology that's going to be in every tool that we use. It's going to be in front of every potential custodian in every litigation. We have to have some guardrails around what that means.

There are so many open questions after these rulings:

- Is AI generated content ESI that is discoverable under Rule 26 and its state equivalents? In the right circumstances, that answer is likely yes, but what are those circumstances? The *OpenAI* decisions seem to say yes.
- What is discoverable? Is it the prompts or the output from the tool? What are the parties' obligations to provide notice to these tools to retain information?
- And what are the tools—the companies who own these tools—obligations to preserve that information?
- What are the circumstances that have to exist for the attorney-client privilege to apply? What about the work product doctrine?
- What about the confidentiality issue and the language of the privacy policies of the AI platforms?

These are all open questions to which we do not have answers. But what we do know is that just as District Judge Rakoff stated in his opinion, more than 80 million people are using these tools. We aren't going to stop that. It's only going to grow. Now, if you've watched the news this week, you've seen that the government is integrating those tools at every step of national defense, national security. So we need guidance from the

courts. We need guidance from the Rules Committee.

What it does mean is this – you have to get in the game on understanding what is out there, how it applies to the facts of your client’s case, and make sure you are taking steps to talk within your firm and your clients about these issues. What happened to the defendant in *Heppner* is likely happening or could happen to just about anyone out there on any kind of litigation. Anything. Talk about process, identify risks in technologies you’re currently using. If your custodians are talking to Microsoft CoPilot, you’re talking to an AI platform. It’s just ChatGPT on the back end. How will the courts handle that?

Conclusion

If you were one of the more than 4500 registrants of the University of Florida eDiscovery Conference last week, I implore you to go and watch the final session of the conference. It’s one you ABSOLUTELY need to see. [Tessa Jacob](#) of Husch Blackwell, [Meribeth Banaschik](#) of EY, [Lea Bays](#) of Robins Gellar, and [United States Magistrate Judge Gary Jones](#) put on an incredibly informative mock hearing about these very issues in which an HR analyst used CoPilot to run numbers for use in a RIF determination. Watch it. Play it out in your head how this will go in your workplace. Think about the millions of other scenarios that could be happening right now. You can follow Minerva26 for all of the updates on the case law and the rulings you need to know about using that Generative AI tag.

That’s our Case of the Week for this week. I’m your host, Kelly Twigger and I invite you to subscribe to our blog at Minerva26 to follow our Case of the Week updates, as well as to subscribe to the Meet and Confer podcast. If you value this content and want us to keep providing it, please share it with your colleagues on social media so that we can continue to reach as many folks as possible and get this word out about the risks, the questions that need to be asked on these topics. If it’s about the discovery of ESI, it’s covered in Minerva26, your discovery strategy platform. Thanks for joining me. See you next time.