

## When Your Expert Uses AI: Prompts as Methodology in *Conservation Law Found., Inc. v. Shell Oil Co.*

### INTRODUCTION

---

Welcome to the Meet and Confer podcast and our Case of the Week segment for this week. My name is Kelly Twigger, and this week we are back — for the third time this Spring — on the topic that has dominated discovery practice in 2026: generative AI. Yes, I know you're probably tired of hearing about it, but this week the question is one we have not yet had a federal court answer. When your expert witness uses AI to develop her report — to filter a document production down to a working subset, to test a hypothesis, or to analyze evidence — are the prompts that she used discoverable under Rule 26?

We finally have an answer, but there are caveats. The answer is yes. And — this is the unusual part — the answer is being delivered to us not in a written opinion, but in a minute order on the docket.

Now I want to flag two things up front before we get into the analysis, because they both shape the strategic value of this ruling. The first is that since the magistrate judge issued the order, the district court has stayed it. On June 3, 2026, United States District Judge Vernon Oliver granted CLF's motion to stay the magistrate's order pending the district court's review of CLF's Rule 72 objection. Shell's response is due tomorrow, June 10, 2026, the day after I'm recording this episode. After that, District Judge Oliver will decide whether to keep the stay in place or whether to uphold or vacate Magistrate Judge Farrish's order. So the ruling we're discussing today is currently on pause. We don't know yet whether it will survive, but the framework is now public, and that's what we're going to talk about.

The second thing I want to flag is that while we're waiting for the court to address the expert AI question, the Florida Supreme Court made its own move on AI accountability. On May 28, 2026, it amended [Rule 2.515 of the Florida Rules of General Practice and Judicial Administration](#).

That amendment takes effect on June 15, 2026 — less than a week from today. We're going to come back to that, because it's part of the same conversation.

Now, as always, I'll provide a link to the docket entry in Minerva26 along with the prior decisions that we'll be discussing today. If you're a Minerva26 subscriber, you can use the Generative AI issue tag to be notified the moment new decisions on this issue drop. They are coming fast and you need to be reading them.

All right, let's take a look at our case from today.

### **The Case: *Conservation Law Found., Inc. v. Shell Oil Co.***

---

The case is [Conservation Law Foundation — or CLF — v. Shell Oil Co.](#) It's proceeding in the United States District Court for the District of Connecticut. And the order we're discussing today was entered on May 18, 2026, by United States Magistrate Judge Thomas Farrish.

Now, before we go any further, I want to be transparent again about what we're working with. **This is a docket text order. There is no written opinion.** It's a minute order entered in text only on the docket. What I'm reading from — what we're looking at today — is the judge's text on the docket at ECF number 970. There's no F. Supp. cite, there are no Westlaw star paragraphs. And if you cite this, you'll have to verify the docket entry through PACER first. Every decision in our Minerva26 database is available publicly, so you can click the link that I give you to view the docket entry yourself, but I want you to understand where it's coming from.

And I want you to sit with that for a second, because it has a strategic implication that we will come back to at the end of our topic discussion today. The AI discovery doctrine is being built case by case, motion by motion, and a significant amount of that building is happening on the docket — not in published opinions. I've seen three docket orders in the last week alone that are in our database and impacting discovery. If your case law research is limited to large databases, you may be missing the doctrine as it is actually being made here. And while the courts may not encourage citing docket entries, the reasoning and rationale behind them can shape your arguments going forward that will show up in published opinions. And as we've seen on the case law on ESI, the rules change regularly. We're constantly citing district courts from other non-precedential jurisdictions. We're constantly looking at the reasoning that other judges are applying because this technology and the changes to the law are happening so quickly. So you may be able to cite those docket orders soon enough.

### **Why this case matters beyond discovery**

---

Now let's talk for a second about why this case matters beyond discovery. I want to take a minute on the underlying case because *Conservation Law Foundation v. Shell* is not just any commercial litigation. It is one of the most consequential environmental cases in the country.

The Conservation Law Foundation, or CLF, filed citizen suits under the Clean Water Act and the Resource Conservation and Recovery Act against Shell and the operators of a bulk storage and fuel terminal on the waterfront in New Haven, Connecticut. The theory of the case is that the defendants failed to design, maintain, modify, or operate their oil terminal to account for the effects of climate change — sea-level rise, more frequent and more severe storms, storm surge — and that those failures create the imminent risk that the next significant weather event will discharge oil and toxic chemicals into the surrounding waters and neighborhoods.

This case, along with CLF's parallel case against Shell in Rhode Island, is among the first in the United States to hold fossil fuel infrastructure operators accountable, not for greenhouse gas emissions, but for failing to prepare their physical facilities for climate change. Per CLF, these are also the first U.S. climate adaptation cases of this kind to reach the discovery phase. That's what CLF argues.

Now that is the context in which CLF retained Dr. Naomi Oreskes as its expert. Dr. Oreskes is a Harvard professor of the history of science whose research has focused on climate disinformation by fossil fuel companies, including in her book *Merchants of Doubt*, which she co-authored. That is the context in which she used AI to work through Shell's document production. So the stakes on the discovery order we are about to discuss are not abstract. They land in one of the most carefully watched environmental cases in the country.

## FACTS

---

Let's talk about the discovery ruling. As always, the timeline is critical in discovery issues. And here's the sequence.

CLF retained Dr. Oreskes as a testifying expert. Her team needed to work through Shell's document production in order to write her report. To make the volume of millions of documents manageable, they used artificial intelligence to cull the document production down to a subset of documents that Dr. Oreskes could actually analyze.

Several months after the report was issued, Shell moved to compel production of Dr. Oreskes' reliance materials — the documents that she relied on in producing her report. Now that's critical: the documents she relied on. That motion came in at ECF order number 941. Per the Court's earlier discovery order at ECF 473, the parties filed simultaneous letter briefs. The Court then ordered them to meet and confer to narrow the dispute, which they did. And by the time the parties came back, they had resolved everything except Shell's request for the AI prompts that Dr. Oreskes used and the outputs those prompts generated. The Court held a hearing on that single remaining issue on May 14, 2026. And four days later, on May 18th, Magistrate Judge Farrish entered the docket order granting Shell's motion.

CLF made three arguments and the Court rejected all of them. Let's walk through them because each one of them is one that you may see or want to make in your case.

## ANALYSIS

---

### Argument one: AI prompts are not within the scope of Rule 26(b)

---

CLF's first position was that AI prompts used by an expert witness are simply not within the scope of discovery at all. The Court rejected that in a sentence. Magistrate Judge Farrish held that an expert's methodology is fair ground for discovery, citing [Macchia v. ADP, Inc.](#), and that under the facts of this case, the process by which Dr. Oreskes culled Shell's production into a subset to be worked with is an aspect of her methodology.

And here's what's critical to understand: the Court is not treating AI prompts as a new category of evidence. It's treating them as part of the expert's method. When an expert uses a tool to filter, rank, summarize, or classify documents, the way she uses that tool is part of the methodology. That principle is not new. We've had discovery of technology-assisted review protocols going back to *Da Silva Moore* in 2012. We've had discovery of search term lists, of sampling protocols, of coding instructions. AI prompts are really just the next iteration of that same idea.

What is new is the application to a testifying expert. And here's where this case fills a gap in the body of law we have been building since last year.

### Where this case fits in the line we have been covering

---

If you have been listening, you know that we have spent the last few months mapping out how courts are going to deal with AI prompts in discovery. We covered multiple decisions in the OpenAI matter. And then the [Heppner](#) and [Gilbarco](#) decisions together in [Episode 183](#). Those were the first two federal rulings on AI privilege. Then we covered the [Morgan](#) case a few weeks later, which is the most thorough civil treatment we have so far. If you haven't listened to those episodes, go back and start there because today's case is the next chapter in that story.

And here's how that spectrum has developed and why *Conservation Law Foundation* matters.

In *Heppner*, the user was a represented criminal defendant — a corporate executive indicted on more than \$150 million in securities fraud — who used Claude on his own initiative without his lawyer's direction to prep defense strategy notes. The FBI seized those notes during a search of his home. District Judge Rakoff held no privilege and no work product in the notes. The defendant had not used the AI at counsel's behest. The platform was not an attorney. The privacy policy of the AI tool permitted disclosure to third parties. And as District Judge Rakoff put

it, "non-privileged communications are not somehow alchemically changed into privileged ones upon being shared with counsel."

In *Warner v. Gilbarco*, decided the same day as the oral decision in *Heppner*, the user was a pro se plaintiff who used ChatGPT for her own analysis of her own case. Magistrate Judge Patti held that her prompts and outputs were not discoverable, and that even if they were, they were protected by the attorney work product doctrine — which a pro se litigant can assert. And Magistrate Judge Patti delivered the line that I think is going to be quoted for years on this issue: that AI platforms like ChatGPT are "tools, not persons." And the theory that disclosure to the tool waives work product would "nullify work-product protection in nearly every modern drafting environment, a result no court has endorsed." So that's a difference between that and *Heppner*, but *Gilbarco* was a civil decision, whereas *Heppner* was a criminal one.

In *Morgan*, Magistrate Judge Braswell extended that analysis to a pro se plaintiff in employment discrimination. The defendant moved to compel disclosure of the AI tool itself and to amend the protective order to require that disclosure. Magistrate Judge Braswell, who, as I noted in our *Morgan* episode, is one of the most informed jurists writing on AI and the law today, held that work product protects pro se AI use. That disclosure to an AI is not disclosure to an adversary, drawing analogies from *Carpenter v. United States* and *Warshak* on third-party privacy. She also gave us the most comprehensive AI-specific protective order language that we have to date.

But what none of those decisions told us is what happens when the person at the keyboard is a testifying expert witness. That was the gap that *Conservation Law Foundation* fills.

Now, each of these decisions that we just talked about are sensitive to who is at the keyboard and why. And the question I want every litigator listening to think about is: where does your client sit on that map?

## **Argument two: the prompts were "notes" covered by a Rule 29 agreement**

---

All right. The second argument that CLF made is one we've not really seen yet, and it's a strategic one worth understanding. CLF argued that the parties had entered into a Rule 29 agreement not to take discovery of each other's "expert notes, drafts, or communications needed by and made during the report drafting process." Counsel for CLF argued at the hearing that Dr. Oreskes' AI prompts qualified as notes within that carve-out.

Magistrate Judge Farrish was not having it. He reminded everyone that he had already addressed Rule 29 in the case and that before a court denies otherwise-relevant discovery based on a Rule 29 stipulation, the agreement "must be quite clear." Calling an AI prompt a note is not quite clear, according to the judge. So that argument failed.

And here's the part that litigators should steal — this is the practical move for you, your firms, and your clients. **Rule 29 agreements have to say what they mean.** If you are negotiating a discovery stipulation today and you intend it to cover AI inputs and outputs, you have to name

them. Do not assume "notes" or "drafts" or "communications" will be read to sweep in AI artifacts. After this case, if it holds up, they will not.

Now this connects directly back to the conversation we've been having since [Episode 183](#). I've been saying in every episode that you need to be talking to your custodians and your experts about their AI use. You need to know what platforms they are using, what they are doing with them, and whether their prompts are preserved. The same conversation now applies to your protective orders and your Rule 29 stipulations. If you want to protect AI artifacts, write them down by name. Use terms like "AI prompts," "AI queries," "AI outputs," "embeddings," and "model parameters." Generic language will not be read to cover them if this decision is what we go by.

And if you are on the receiving side of an AI prompt request and you want to protect it, your strongest move is not to retroactively argue the prompt was a note without further clarification of what a note means. Your strongest move is to negotiate the carve-out in writing up front. This is another situation where you have to be thinking early about these kinds of situations, planning and preparing for them. The after-the-fact remedies are going to catch up with you. Magistrate Judge Braswell's protective order language in *Morgan* is one starting point, but it does not address experts. So you're going to need to draft your own language there. Think about the Rule 29 stipulation. It's likely on the docket in the Shell case — you can look at it — but again, it needs to be modified. The judge held here that the language counsel used did not cover the AI prompts.

## The Technology at Issue

---

Now, before we get to CLF's third argument, it's important to understand the technology that was at issue here — the technology that Dr. Oreskes and her assistant used to filter documents.

CLF disclosed that Dr. Oreskes and her assistant used OpenAI's GPT-4o and GPT-4.5 models accessed through the Microsoft Azure OpenAI Service via the secure Azure API. That is the enterprise deployment of OpenAI's models — not the public ChatGPT consumer product. The documents were processed on a Microsoft Azure cloud server that CLF says was rented for exclusive use on this litigation.

And that distinction matters. It also tells you you can do it if you need to. It matters because Shell's original motion had a separate piece challenging whether uploading confidential documents to a third-party AI platform violated the Standing Protective Order. And that piece of the dispute was actually resolved at the parties' four-hour court-ordered meet-and-confer in April. By the time the May 14th hearing happened, the Azure configuration was no longer in dispute. The only thing left was the prompts and the outputs themselves. I'm going to guess a lot of that had to do with Magistrate Judge Braswell's *Morgan* opinion, in which she said you have three requirements that you have to meet in order to put confidential information into an AI

tool, and the enterprise deployment of Azure met those requirements. That's not coming straight from the opinions — that's coming straight from my mouth.

The broader point is this: the tool an expert uses is itself a question you need to be asking. The risk profile of GPT-4o running on an enterprise Azure instance with no model training and no data retention is fundamentally different from the risk profile of someone typing the same prompts into the consumer version of ChatGPT. We've been on that point all year, and we just discussed how that comes from Magistrate Judge Braswell's protective order language in *Morgan. Conservation Law Foundation* is the reminder that the line matters at the expert engagement stage too. Have those conversations with your expert at the outset. Have a preservation policy and plan in place.

### **Argument three: there is nothing to produce**

---

All right, now let's turn to CLF's third argument — that Dr. Oreskes did not use prompts at all. She used search terms. And CLF said that it had already produced the search terms.

And that's not just lawyer framing. CLF anchored the argument that Dr. Oreskes used search terms in the AI tool rather than prompts in Dr. Oreskes' own deposition testimony. And that testimony is worth reading because it shows you exactly how AI methodology disputes are going to play out. When Shell's counsel asked Dr. Oreskes about the prompts she had used, Shell's counsel framed a prompt as something different from a search term. A prompt in counsel's framing implied generative analysis, where the AI is doing the thinking. Dr. Oreskes pushed back. She told the lawyer on the record that what she and her assistant did was not generative analysis at all. The AI sifted documents by keyword. Then she read every document the AI surfaced. Nothing in her report was generated by AI. The prompts Shell is asking for, CLF argues, are nothing more than search terms, and CLF has already produced those.

Now, you might be wondering, as I was, what is the difference between a prompt and a search term anyway? Both of them are just text that you type into an AI tool. And the honest answer is that at the technical level, there is no difference. Anything you send to generative AI through an API is processed by the model as a prompt. There is no "search term mode" hidden somewhere inside ChatGPT. Whether you type "sea level rise" into Google or send the identical text to GPT-4o, those are not the same operation. Google runs a keyword index. GPT-4o is a generative language model. It is interpreting what you mean, pulling in concepts that share meaning but not literal text, and ranking documents by inferred relevance — not literal word match.

So when Dr. Oreskes' team called what they did search terms, they were doing legal work with that label. They were trying to make the analogy to a Google search or to old-school keyword culling, which has never been discoverable in this way. But the analogy breaks down at the technical level. The way you phrase a prompt to a generative AI changes what the prompt

surfaces. Different phrasings produce different document sets. Different document sets produce different opinions. The prompt is not plumbing. The prompt is methodology.

And that's likely why Magistrate Judge Farrish ruled exactly the way he did. He did not have to take a position on the philosophical question of whether the inputs were search terms or prompts. He just had to look at what the technology actually was. The technology was a generative model. The inputs were prompts. The prompts shaped the methodology. Methodology is discoverable. And that was the end of the inquiry. So we'll see how the district court handles that on the objection.

Here's the part that litigators need to internalize. If you're going to use generative AI in connection with expert work, you cannot relabel your way out of discoverability. Calling a prompt a search term does not change what it is. The court is going to look at the actual technology, how it functions. And so will opposing counsel, and so will the jury.

But here's my question — if CLF disclosed the search terms, what is Shell actually fighting over? And the honest answer is that "search terms" is a much narrower thing than the actual methodology. And CLF's framing is doing some legal work to make it sound complete when it's likely not.

Here's what's missing. When you use a generative AI to filter documents, you don't just send the model a keyword. You send the model an instruction. A prompt is an instruction: review this document, tell me if it discusses sea-level rise, output a list of document IDs. The keyword is one variable inside that instruction. But the instruction itself — the wrapper prompt — is methodology, and CLF has not disclosed it.

Second, AI filtering is iterative. You run a query, you see the results, you refine, you run it again. CLF's own evidence acknowledges that the prompts and outputs were used interactively. That iteration shapes which documents end up in the working set. And CLF did not preserve any of it.

Third, even if you know which keyword went in, you don't know which documents came out. The AI returns an ordered list — its judgment about relevance. That ranking is methodology that Dr. Oreskes effectively delegated to the model. Shell does not have that ranking to be able to question the expert's decisions.

So when CLF says it disclosed the search terms, it's disclosing one variable from a multi-step process: the wrapper prompt, the iterative refinements, and the model's intermediate rankings — none of that has been produced. That's the gap that Shell is trying to recover here. And it's the same gap that Magistrate Judge Farrish was looking at when he ordered production.

So the takeaway for litigators is this: if you're using generative AI to filter documents for an expert, the disclosure obligation is broader than a list of keywords. It is the entire instruction set — the wrapper prompts, the refinements, and the outputs. If you want to argue that you've disclosed everything, you need to be able to show you've preserved everything. Disclosing

search terms is a partial disclosure of a multi-step process. And partial disclosure is the gap that opposing counsel and the court will both find.

Now, the Court walked through its own analytical framework on the "nothing to produce" piece. And this is worth understanding because it's a frame that could be used against you. Ordinarily, when the producing party tells the court "we don't have it," and the court has no reason to doubt that representation, the court will not order production of the impossible. But that good faith rule does not apply when the requesting party has a strong reason for disbelieving the representation — a reason backed by solid evidence, not mere suspicion. And that comes from the *Lewis v. Doe* opinion, which is an opinion that Magistrate Judge Farrish himself authored.

So here Shell had that solid evidence that was needed. Dr. Oreskes' assistant had submitted a declaration that used the word "prompt." And once that word was on the page, Shell had a non-speculative reason to believe that prompts existed. And that meets the standard for the court. CLF's blanket denial that the team had used the word "prompts" was not going to carry the day.

So the Court did something kind of elegant. It ordered CLF to revise its responses to Rule 33 interrogatories and any Rule 34 requests for production that called for the disclosure of AI prompts or queries used by Dr. Oreskes or her team. If after a diligent search CLF really has nothing more to produce, fine — but it has to say so under oath, signed by the appropriate person under Rule 33 or 34. And Rule 37(b) sanctions become available if that representation later proves untrue. CLF was given until June 1, 2026 to serve the revised responses. Remember that Rule 37(b) is the violation of court order, right? So the court ordered it. If you violate that order, you've got sanctions under Rule 37(b). That's not Rule 37(e), where we're on failure to preserve.

Now that's the trapdoor. The Court is not prejudging whether prompts exist — it's forcing CLF to put its credibility on the record. If prompts surface later in deposition, in metadata, in the AI vendor's logs, then the Court has the sanctions hook ready.

Now, looking back at the record in this case, CLF has said that this information doesn't exist multiple times. So I'm not sure that's going to change. That's the natural question: do the prompts actually exist? They've said repeatedly under oath, in signed objections, in counsel statements at the May 14th hearing, and in the Rule 72 objection that they do not exist. The trajectory of the representation has tightened over time — from "no prompts were stored" to "no complete log was preserved" to the current position that no log exists in any form.

But here is what the Court actually decided. Magistrate Judge Farrish did not find that prompts exist. He found that Shell had a non-speculative reason to disbelieve CLF's representation, and that CLF therefore had to put its representation on the record under Rule 33 or 34, signed under oath with Rule 37(b) sanctions available if anything later turns up. That is trust but verify. It is not a finding that CLF is hiding the ball.

And here's the deeper exposure that CLF is sitting on, which really the briefing doesn't fully grapple with. Even if the assistant's local system retained nothing about the prompts, the data that CLF processed through Microsoft Azure OpenAI Service is presumptively within CLF's possession, custody, or control — because CLF, through the assistant, was Microsoft's customer. The producing party cannot stop at its own laptop. It has to make a reasonable inquiry into where its data lives. And that includes asking the cloud vendor. So if CLF wants to stand on "nothing exists," they may need to be able to show that they asked Microsoft and Microsoft retained nothing.

And if Microsoft did retain something — Azure's own documentation acknowledges that the Azure OpenAI Service may retain prompts and outputs temporarily for abuse monitoring — then CLF has to produce it. It's been months now since this request happened, since these letter briefs have been filed. It's possible — maybe not likely — that abuse monitoring has still captured some things. And if Microsoft retained nothing because the deployment was configured for zero retention, then CLF is sitting on a Rule 37(e) preservation question.

We talked about this in our [PharmacyChecker episode](#). When a source of ESI is going to disappear — whether it is Universal Analytics, as it was in that case, sunsetting, or an Azure session ending — counsel cannot delegate the preservation decision to a technical team member and walk away. The duty to take reasonable steps to preserve is a duty that counsel has to supervise. The assistant here was a research assistant. He is the person who decided not to export a complete log. Was CLF's litigation counsel involved in that decision? The briefing doesn't say. That may be the exposure.

So when I look at where this case actually lands, the "nothing to produce" question is not resolved. It's just deferred. Either the prompts exist on CLF's systems and CLF will have to produce them, or prompts exist on Microsoft's systems and CLF will have to retrieve them. Or nothing exists anywhere and CLF is in Rule 37(e) territory for failure to preserve a work stream it knew was central to expert testimony. None of those paths is a clean win for CLF.

## **What CLF is actually arguing on the stay — and why it matters strategically**

Now let's talk about what CLF put in its Rule 72 Objection, which caused District Judge Oliver to stay Magistrate Judge Farrish's decision. And that briefing is publicly available. And the arguments that CLF is making to the district court are more interesting, more developed, and more useful to litigators than what made it into Magistrate Judge Farrish's docket text order. If you're dealing with this issue in your own case, the CLF brief at ECF 976 is worth reading in full. There are six grounds listed there, and I want to talk about three of them.

The first one is a procedural defect. CLF argues that Shell never actually served Rule 33 interrogatories or Rule 34 requests for production for the AI prompts. What Shell served was a series of follow-up emails and meet-and-confer letters. CLF voluntarily responded to those —

and that's why we have the signed September 25, 2025 objections. But the objections were courtesy responses, not responses to formal discovery requests.

And here's the doctrine: an informal email request cannot be enforced through a Rule 37 motion to compel. CLF cites several decisions, *Wells Fargo* and *Schwartz*, in which it says plainly, "the entire enforcement mechanism of Rule 37 contemplates the parties having formally resorted to the underlying discovery rule, in this case Rule 34, rather than a casual, informal request contained in a letter." Then CLF goes a step further and argues *Wynne v. Town of East Hartford*, which holds that Rules 33 and 34 are not even the proper vehicle for expert discovery in the first place — that depositions and Rule 45 subpoenas to the expert are proper.

Now read those together. CLF is arguing that Magistrate Judge Farrish ordered it to revise responses to Rule 33 and 34 requests for production that were never served, and authorized Rule 37(b) sanctions for any inconsistency in those responses, when the underlying enforcement mechanism does not exist. That's the kind of clean procedural argument that a district court can use to vacate an order without ever reaching the merits of whether AI prompts are discoverable expert methodology. Now, I don't know where we are in the case management order or in terms of discovery process whether Shell can at this point go back and provide those actual discovery requests in terms of Rule 33 and 34 to trigger that obligation under Magistrate Judge Farrish's order.

The second argument that CLF makes on its objection is this sword-and-shield framing. Shell has a pending Daubert motion to exclude Dr. Oreskes' testimony. And the basis for the Daubert motion is, in part, the same alleged disclosure failure on AI prompts that drove the motion to compel. CLF argues that Shell is using the same alleged gap as both sword — the Daubert exclusion if the prompts are not produced — and shield as the motion to compel to force production. And you can't have it both ways. And the timing really supports CLF here. Shell's request came 88 days after Dr. Oreskes' expert report was disclosed, when the CMO's guideline for motions to compel was 30 days. Both arguments are in front of District Judge Oliver right now.

The third argument that CLF makes I want to cover is the exchange at the hearing. There's a moment from the May 14 hearing that I think every litigator should pay attention to because it shows how courts are wrestling with this. Magistrate Judge Farrish pressed Shell's counsel on why it actually matters which documents reached Dr. Oreskes. And Shell's answer was this: "If there are documents in Bates Nos. 1 million to 2 million that contain things that are directly opposite of the opinion she's offering, and she didn't even consider them, she didn't even give herself the chance to consider them because they're not in the set, that's a pretty potent argument for our Rule 702 motion. And failing that, it is a really powerful argument for the jury."

Shell is being honest with the Court here. The purpose of the prompt discovery is not just to test methodology — it's to find the documents that Dr. Oreskes did not consider and to argue that she should have. And CLF's response is that Rule 26(a)(2)(B)(ii) does not require disclosure of documents that an expert did not consider. That's key. The expert disclosure obligation is

forward-looking. It covers what she relied on, not a roadmap of everything she could have looked at and didn't. And that's another doctrinal question. It is the kind of question that doesn't get answered in a docket text order, and it didn't get answered here.

## **Bringing in Florida — Rule 2.515(d)(2)**

---

Now let me take you to the second piece of news that I flagged at the top, because it lands at the exact same moment here and in the same conversation on the same underlying principle.

On May 28, 2026, the Florida Supreme Court amended Rule 2.515(d)(2) of the Florida Rules of General Practice and Judicial Administration. And that amendment takes effect in six days on June 15, 2026.

Here's what the amended rule does. By signing a court filing in Florida, every signer now represents that the legal authorities identified exist and are accurately cited. The amendment expressly authorizes courts to impose sanctions for any filing inconsistent with that representation — reprimand, contempt, striking the document, dismissal, costs, or attorney's fees — after notice and an opportunity to be heard. The court issued a companion administrative order that preempts the patchwork of circuit-level AI disclosure orders that had grown up around the state. There's now one statewide standard to be applied.

Now I want to be precise about what that rule does and does not do. It does not require disclosure of AI use. We've seen a lot of judicial orders requiring disclosure of AI use. It does not require certification that AI was not used. It does not ban AI. What it does is put accountability on the signer. If a hallucinated citation lands in a brief, the question is no longer whether the AI did it. The question is whether the signer verified the authority before filing. The signer is on the hook.

Now, why am I bringing a Florida court filing rule into a Connecticut expert witness case? Valid question. And the answer is really because the through-line is the same. And it's the same through-line that District Judge Xavier Rodriguez was pointing at when he raised the rethinking-privilege-in-the-age-of-AI question at UF that we've been talking about on all these episodes. **Courts are converging on a single principle: AI does not change accountability.** The lawyer is still responsible for what is filed. The expert is still responsible for her methodology. The litigant is still responsible for the responses she serves under oath. The fact that a tool produced the output does not insulate any of them.

For the expert side, the downstream effect of Florida's new rule is significant. If your expert is preparing a report in a Florida-venued case and her draft will eventually be incorporated into a motion, a Daubert brief, or any other court filing, the lawyer who signs that filing is making a representation about every authority cited in it — including authorities the expert pulled with an AI tool. Accountability does not stop at the expert's office. It runs all the way to the signature block.

And this is the part that's going to matter for your firms. Our procedures cannot just be about lawyers and AI. They have to extend to the people whose work product gets folded into our filings. That includes experts, and it includes the consultants and assistants that those experts use.

## The docket-order issue — and why it matters strategically

---

Now let's talk a little bit about the docket order here and why it matters strategically. Let's come back to the fact that *CLF v. Shell* is a docket text order, because I said at the top that this matters strategically. And let me explain that a little bit.

A docket text order has the full force of the court's authority in the case that it was issued in only. But it does not have the same precedential reach as a written, published opinion. There's no headnote, there's no syllabus, there's no F. Supp. version with paragraph numbers. You can cite the ECF entry, but the judge reading your brief next year may not pull the docket to check it. And the holding lives or dies on the text the judge typed into the docket — that's exactly the text that I read to you earlier. So that holding lives or dies on just that text that's in the docket. If you're a Minerva26 user, you can use the publicly available opinion from Minerva26 to cite to it.

Now, that is layered on top of the stay that we discussed at the start of this episode. District Judge Oliver has paused the order while he reviews CLF's Rule 72 objection. And until he rules, the magistrate's order is not enforceable. So the precedential value of what we have today is doubly limited. It is one magistrate judge in one case in one district in a docket text order that has been stayed pending district court review.

And that cuts two ways. Don't over-weight this decision in your current briefing — there is still a meaningful chance that District Judge Oliver vacates or modifies it. But don't under-weight it either. The reasoning is on record. Other magistrate judges and district judges across the country are watching this. The next time the expert AI question comes up — and it will — opposing counsel will know about this order, and so will the judge.

The bigger lesson here is one that I have been hammering for five years on this podcast, and it has never been more true: **the discovery doctrine is moving faster than the published case law can keep up with**. If your research practice is limited to what makes it into the books, you are missing the doctrine that is actually being made.

## TAKEAWAYS

---

All right, let's talk about takeaways. And here is what I see are counsel's obligations today in light of this ruling.

First, if your testifying expert is using AI in any way to develop her opinions, treat the prompts as discoverable methodology now. Preserve them in real time. You can try to protect them later, but

do not let them live in a chat history that disappears when the browser closes or is not subject to a retention policy. This is a conversation you need to have with your expert at the engagement letter stage, not after a motion to compel lands. We'll see if the district court's ruling changes any of that, and if it is a written opinion and not just a minute entry on the docket.

Second, build the AI methodology into the expert report itself. If prompts are going to be produced, you want them produced with the rationale — the limiting instructions, the validation steps, the documentation of how the output was used. A prompt produced cold two months after the fact is a gift to a cross-examiner. A prompt produced inside the methodology section of the report is a defense to how it was used.

Third, if you want to protect AI artifacts under a Rule 29 stipulation or in a protective order, name them. Notes, drafts, and communications will not be read to sweep them in. After *CLF v. Shell*, you will have to use the words: AI prompts, AI queries, AI outputs, embeddings — whatever is appropriate.

Fourth, when you are on the receiving end of "we have nothing more to produce" in any AI dispute, look for the credibility hook the way Shell did. In *Conservation Law Foundation*, the hook was a single word — "prompt" — in a member of the expert team's declaration. Read everything. The credibility evidence is usually already in the record. We've talked about this before in the PharmacyChecker case and in other episodes on the Case of the Week. You have to have that hook to get over the assumption that the information doesn't exist.

Fifth, if you practice in Florida, or any of your matters are venued there, calendar June 15, 2026. After that date, the Rule 2.515(d)(2) certification is operative on every single filing. Make sure your firm has a written verification protocol for every authority cited in a brief, regardless of whether AI was used in drafting. The rule does not require disclosure — it requires accuracy. The way you protect yourself is process.

Sixth — and I cannot say this enough — start treating the four cases we have now covered as a single framework. *Heppner* tells you what happens when the user is operating outside counsel's direction on a public platform. *Warner* and *Morgan* tell you what happens with pro se litigants. I had a client call me last week and ask me about how they should be disclosing the use of AI to their clients, how that information has to be included. Another client asked me what needs to be included in their engagement letter to their clients about use of AI and about confidential information being put into AI. All of these things are being formed by this case law that we're seeing. Knowing where on the map your case sits is how you decide what to protect and what to produce and what to negotiate.

And one final point, because I want to get this in front of every white-collar lawyer listening. We talked about this in the *Morgan* episode. *Heppner* did not just apply to one defendant — it applies to every represented party who uses AI without their lawyer's direction. If you represent a client — a corporate executive, the target of an investigation, anyone — who is using AI on

their own to think through their case, that is a conversation you need to have today, before the materials become an issue.

## WHAT HAPPENS NEXT — AND CONCLUSION

---

All right. Now let me close with the prediction I owe you from the top of the episode. The honest answer is that District Judge Oliver has several paths in front of him. And I do not think this is a one-outcome case.

The narrow path that he may take is procedural. District Judge Oliver could vacate the order on the Rule 33/34/37 procedural defect alone — no formal requests served, so no enforceable Rule 37 motion, so the magistrate's order rests on a discovery mechanism that does not exist. That outcome leaves the substantive question — whether AI prompts used by a testifying expert are discoverable as methodology — completely unanswered. The framework that has been getting all the press coverage could come completely off the books. And the next time this question arises, the next court could be writing on a blank page, potentially with this as a guide.

The broader path is substantive. District Judge Oliver could affirm that an expert's AI methodology is fair game under Rule 26, that prompts are a part of that methodology, and leave Magistrate Judge Farrish's framework in place, but adjust the enforcement piece given CLF's representation that no prompt logs were preserved. That outcome keeps the doctrine alive, but takes some teeth out of the order.

And then there are the paths in between. District Judge Oliver could split the difference — affirm the methodology principle, vacate the Rule 33/34 enforcement mechanism, and direct the parties to a different procedural posture. Or he could remand for further factual development. Or he could wait on the related Daubert motion, which is in front of him on the same record.

The point is, the framework that everyone is treating as settled is still very much up for grabs. I'm not going to call it. What I am going to do is watch the docket, and when District Judge Oliver rules, we'll come back to this on a future episode and walk through what the district court did on each of those arguments. Whatever happens next here is going to be a teaching moment for the body of law that we have been building all Spring.

That's our Case of the Week for this week. We'll have a link to the docket entry and whatever comes next from District Judge Oliver in *Minerva26*, along with *Heppner*, *Warner*, *Morgan*, and the Florida Supreme Court order. If you're a *Minerva26* subscriber, again, use that Generative AI issue tag to be notified the moment a new decision drops on these issues. They are coming fast and furious.

If you have questions or you want to dig deeper, bring them to the next Meet and Confer session or drop them in the comments. And please share this episode with the litigators in your network who work with expert witnesses or who have any matter where AI is being used to

process documents. Staying current on this body of law is critical. It is timeliness, and it is hard to do on your own.

I'm your host, Kelly Twigger. I invite you to subscribe to our blog at [Minerva26.com](http://Minerva26.com) and to follow all of our Case of the Week updates and to subscribe to our Meet and Confer podcast wherever you get your podcasts. If it's about the discovery of ESI, it's covered in Minerva26, your discovery intelligence platform. Thanks so much for joining me, and I'll see you next time.