



## Court Orders Disclosure of Expert Witness's AI Prompts: What Litigators Need to Know

A federal magistrate judge in Connecticut has issued what appears to be the first decision of its kind: an order requiring a party to produce the generative-AI prompts its expert witness used while preparing her report. The ruling is short, but its implications are not. For litigators, it signals that the inputs an expert feeds to an AI tool – the prompts, the queries, the iterative refinements – may be fair game in discovery as part of the expert's methodology. That is a development with immediate consequences for how experts are retained, how they work, and how they are prepared for deposition and cross-examination. Counsel who assume that an expert's "behind-the-scenes" use of AI will stay behind the scenes should reconsider that assumption now.

### The Ruling

On May 18, 2026, Magistrate Judge Thomas O. Farrish of the U.S. District Court for the District of Connecticut ordered the plaintiff in *Conservation Law Foundation, Inc. v. Shell Oil Co., et al*<sup>1</sup> to produce the generative-AI prompts that its expert, Dr. Naomi Oreskes, used in connection with her expert report. The underlying case concerns a bulk fuel terminal in New Haven and allegations that the facility was not adequately prepared for climate-related risks such as severe weather and flooding.

In preparing her report, Dr. Oreskes, a climate science historian, and a research assistant used a commercially available generative-AI tool to filter and identify potentially relevant documents from the defendants' productions. Oreskes disclosed her use of AI in the report served in May 2025, identifying certain search terms and frequencies. The defendants wanted more. Through her deposition and a series of informal requests, they sought additional information about the prompts and the

system's outputs. When meet-and-confer efforts failed, they moved to compel – and the court granted the motion.

The reasoning was straightforward. An expert witness's methodology is fair ground for discovery under Rule 26, the court reasoned; the process by which Dr. Oreskes culled the defendants' document production down to a workable subset was an aspect of that methodology; and the prompts that drove that process were therefore within scope. The court treated the AI prompts as no different in principle from the formulas, code, or other methodological steps an expert is routinely required to disclose. It rejected the plaintiff's argument that a stipulation protecting an expert's "notes, drafts, or communications" shielded the prompts. Notably, the order has been stayed pending the district court's resolution of the plaintiff's objection, so the last word has not been written.

### **Why the Prompts Were Treated as Methodology**

The central conceptual move in the decision is its characterization of AI prompts as methodology rather than as protected work product. That framing matters because it determines which side of several well-worn discovery lines the prompts fall on.

Under Rule 26(a)(2) and Rule 26(b), the facts, data, and methods an expert relies on in forming an opinion are generally discoverable, even as the 2010 amendments to Rule 26 extended work-product protection to most draft reports and many attorney-expert communications. The plaintiff leaned on that protective zone, arguing that the prompts were used only to "cull a large document universe" and were "never considered by the witness in forming her opinions." The court was unpersuaded. If the prompts shaped which documents the expert ever saw, then they shaped the evidentiary foundation of the opinion – and a party cannot meaningfully test the reliability of that opinion without understanding how the universe of considered material was narrowed.

That reliability concern ties the ruling directly to *Daubert* and Federal Rule of Evidence 702. Generative-AI tools raise pointed questions about reproducibility: the same prompt can yield different outputs, models are updated and retired, and the path from raw data to filtered subset can be opaque even to the user. If an opposing party cannot see what went into the system, it cannot probe whether the

approach was reliable or whether responsive, harmful, or simply important documents were filtered out before the expert ever laid eyes on them. Viewed that way, the prompts are not a peripheral curiosity; they are the audit trail for a methodological choice that affected the opinion's foundation.

It is worth noting what the court did not decide. It did not squarely resolve whether the prompts were protected by the attorney-client privilege or the work-product doctrine, and it confronted a relatively narrow use case – prompts used to identify documents rather than prompts used to draft analysis or generate substantive conclusions. The reasoning, however, will not stay confined to document culling. If methodology is the touchstone, then prompts used to summarize literature, run calculations, or stress-test hypotheses are vulnerable to the same logic.

### **The Broader Discovery Problem**

The ruling exposes a practical problem that most litigation teams have not yet solved: AI prompts and their outputs are frequently not preserved in the ordinary course. A request to produce prompts can quickly expand into demands for full prompt histories, iterative refinements, system instructions, and chat logs – materials that may never have been retained, and whose absence can itself become a point of contention or a spoliation argument. An expert who treated an AI session as a disposable scratchpad may find that the lack of a record is as damaging as an unfavorable record would have been.

There is also a reliability-versus-privilege tension that this decision sharpens but does not resolve. Prompts can blur the line between fact and lawyering. A prompt may contain an expert's neutral query, or it may embed counsel's theory of the case, mental impressions, or strategic framing. Whether a given prompt is a discoverable fact, protected work product, or a privilege waiver is a question courts are only beginning to work through, and the answers will be fact specific. What is already clear is that litigators can no longer treat an expert's AI use as invisible plumbing.

### **Practical Steps for Litigators**

The prudent response is to assume that an expert's AI inputs may have to be disclosed, and to manage the expert engagement accordingly from day one.

Several measures follow directly from the ruling.

First, ask experts about AI use at the outset and revisit the question throughout the engagement. Counsel should know, before a report is served, exactly which tools an expert used, for what tasks, and at what stage. Surprises discovered at deposition are far more costly than disclosures managed proactively.

Second, build preservation into the workflow. If an expert is going to use generative AI, the prompts, refinements, system settings, and outputs should be captured and retained contemporaneously. A defensible record of how the tool was used is far better than a gap that invites accusations of concealment or spoliation.

Third, distinguish carefully between AI used to organize or locate information and AI used to form opinions. The *Oreskes* prompts were used to narrow a document set, and the court still ordered production. Using AI deeper in the analytical process only heightens the discoverability and reliability exposure and should be approached with corresponding caution and documentation.

Fourth, draft disclosures with precision and candor. The expert in this case disclosed AI use but in terms the defendants found incomplete, which fueled the motion to compel. A clear, specific account of how a tool was used – described in the report itself – reduces the room for an opponent to argue that something is being hidden.

Fifth, be deliberate about prompt content. Because prompts can embed attorney strategy and mental impressions, counsel and experts should be conscious that a prompt may one day be read by the other side. That is not a reason to avoid AI, but it is a reason to keep work-product-laden instructions out of prompts that may be produced, and to think through privilege implications before, not after, the tool is used.

Sixth, prepare to litigate the issue both offensively and defensively. When you sponsor the expert, anticipate a motion to compel and have a preservation story ready. When you are opposing an AI-assisted expert, consider early, targeted discovery into the tools, prompts, and outputs as a route to *Daubert* and reliability challenges.

## The Takeaway

*Conservation Law Foundation v. Shell* is a single magistrate-judge order, currently stayed, arising from a narrow set of facts. It is not the last word, and its ultimate fate on objection remains to be seen. But its core logic – that the way an expert uses AI is part of the expert’s methodology, and methodology is discoverable – is intuitive enough that other courts are likely to find it persuasive. As generative AI becomes a routine part of expert work, the question is no longer whether opposing counsel will ask about it but how prepared your side will be when they do. Litigators who treat AI inputs as discoverable from the start, preserve the record, and manage prompt content with care will be the ones who turn this emerging issue from a liability into a controlled, defensible part of expert practice.

*This article was written by [Arnold D. Lee](#), an attorney in the Phoenix, Arizona office of Spencer Fane. For more information, visit [spencerfane.com](http://spencerfane.com).*

-----

1

(Case No. 3:21-cv-00933, D. Conn.) (ECF No. 970)

Click [here](#) to subscribe to Spencer Fane communications to ensure you receive timely updates like this directly in your inbox.