

PUBLICATION

A Legal Framework for the Discoverability of AI

Authors: Edward D. Lanquist, Jr, Nicole L. Imhof, Andrew Jacob Droke
April 07, 2026

A growing body of case law is beginning to define when artificial intelligence (AI)-generated materials may be protected and when they will not. In addition to *United States v. Heppner*, three cases decided before and alongside *Heppner* – *Tremblay v. OpenAI, Inc.*, *Warner v. Gilbarco*, and *In re OpenAI, Inc. Copyright Infringement Litigation* - illuminate the key variables that courts are examining. Collectively, these cases underscore that the identity of the user – client or attorney – and the degree of direction exercised by counsel are often dispositive.

In this alert, we expand upon our [analysis of *Heppner*](#) to provide more context in each of the three aforementioned cases – *Tremblay*, *Gilbarco*, and *In re OpenAI* – and provide a framework for protecting your legal interests as it relates to the discoverability of prompts and AI-generated content.

Attorney Prompts as Work Product: *Tremblay v. OpenAI, Inc.* (N.D. Cal. 2024)

In *Tremblay*, the court confronted the discoverability of AI prompts from the opposite direction: it was OpenAI, a defendant, that sought to discover the prompts and outputs that the plaintiffs' attorneys had generated using the AI tool during their pre-suit investigation.

Judge Araceli Martinez-Olguin held that those prompts were protected work product. The court found that the queries had been "crafted by counsel" as part of a deliberate litigation strategy and contained the "mental impressions and opinions" of the attorneys – the precise content that the work-product doctrine is designed to protect.

The contrast with *Heppner* is instructive and should be internalized by every client. In *Tremblay*, attorneys were the authors of the prompts, and the prompts reflected counsel's judgment about what facts and arguments were relevant to the case. In *Heppner*, the client was the author, acting entirely on his own initiative. The difference in outcome tracks the difference in who held the pen.

The work-product doctrine exists to protect the thought processes and strategies of lawyers preparing for litigation – not the independent research activities of their clients. When a client sits down at a consumer AI platform without any direction from counsel, the resulting documents look nothing like attorney work product, regardless of whether the client later shares them with a lawyer.

Civil Work Product and the "Tool, Not Person" Principle: *Warner v. Gilbarco* (E.D. Mich. 2026)

On the same day *Heppner* was decided, Magistrate Judge Anthony P. Patti of the Eastern District of Michigan addressed a related question in a civil context in *Warner*. The defendant sought production of "all documents and information" concerning the plaintiff's use of AI tools in preparing legal materials. Judge Patti denied the request.

The court concluded that using AI tools to draft legal materials is comparable to traditional work-product activities and rejected the argument that the use of generative AI tools automatically waives work-product protection. Critically, the court held that "ChatGPT (and other generative AI programs) are tools, not persons," and that "no cited case orders the production of what Defendants seek here: a litigant's internal mental

impressions reformatted through software." The court also emphasized that waiver of work-product protection requires disclosure to an adversary or conduct that makes it likely materials will reach adversarial hands, and found that mere use of an AI drafting tool did not meet that standard.

Warner appears, at first glance, to point in a different direction from *Heppner*. The reconciliation lies in context and scope. *Warner* involved a broad discovery request for all AI-related materials in a civil case; the court declined to order wholesale production of an entire category of litigation preparation activities. *Heppner* involved specific documents created by a client who had no attorney direction, using a platform whose own terms disclaimed confidentiality, and who later tried to retroactively assert privilege by transmitting the documents to counsel.

The two rulings are not in conflict: *Warner* stands for the proposition that AI-assisted legal work is not categorically unprotected; *Heppner* stands for the proposition that client-generated AI documents created outside any attorney-client relationship and on a platform with no confidentiality protections will not be shielded after the fact. The critical factors remain attorney involvement, direction by counsel, and the confidentiality posture of the platform.

No Privacy Expectation in Voluntarily Disclosed Prompts: *In re OpenAI, Inc. Copyright Infringement Litigation* (S.D.N.Y. 2026)

Judge Rakoff's written opinion in *Heppner* cited *In re OpenAI, Inc. Copyright Infringement Litigation*, the consolidated multidistrict litigation (MDL) pending before Judge Sidney Stein in the Southern District of New York. This case raised the broader proposition that users do not retain substantial privacy interests in "conversations with [a] publicly accessible AI platform" that are voluntarily disclosed to and retained by the platform. In that litigation, Judge Stein compelled OpenAI to produce 20 million anonymized ChatGPT conversation logs in response to discovery requests by copyright plaintiffs.

While acknowledging that ChatGPT users hold "sincere" privacy interests in their conversations, the court found those interests adequately addressed by de-identification protocols and the existing protective order. The dispositive factor was voluntary disclosure: ChatGPT users, unlike subjects of wiretaps, "voluntarily disclosed" their communications to OpenAI, a distinction the court found fatal to OpenAI's privacy objection.

The OpenAI MDL ruling reinforces what *Heppner* holds at the individual level: when you submit prompts to a consumer AI platform, you are voluntarily sharing that content with a third party, and courts will treat that disclosure as undermining both confidentiality and privilege claims. The ruling also signals that AI companies themselves cannot serve as a barrier against discovery of user data when they have retained that data pursuant to their own terms of service.

What the Cases Tell Us Together: A Framework

Taken together, these cases begin to sketch the contours of when AI-generated materials will and will not be subject to privilege or protected by the work product doctrine.

Protection is most likely when:

- Prompts are crafted by attorneys as part of a deliberate litigation strategy;
- Counsel has directed the AI research;
- Platform terms include genuine confidentiality protections; and
- Materials are never shared outside the attorney-client relationship.

Protection is least likely – and discoverability is most certain – when:

- Client creates the materials independently without attorney direction,
- The platform's terms permit disclosure to government authorities or third parties, and
- The client later attempts to retroactively shield the materials by transmitting them to counsel.

The line courts are drawing is not about whether AI was used; it is about who used it, under whose direction, and with what expectation of confidentiality.

For further analysis tailored to your sector and compliance footprint, please contact the authors – [Edward D. Lanquist](#), [Nicole Imhof](#), [Andrew J. Droke](#) – or another member of Baker Donelson's [AI Team](#).