



Courts Begin Shielding AI Prompts and Outputs From Discovery: New York and Texas Recognize Work-Product Protection for AI-Assisted Litigation Work

Every litigator now confronts a question that did not exist five years ago: when a party uses a generative artificial intelligence (AI) tool to help prepare its case – drafting motions, testing arguments, organizing facts – are the prompts it typed and the outputs it received discoverable by the other side? For opponents, an adversary's AI chat log is a tantalizing target: a candid, time-stamped record of theories considered and abandoned, weaknesses acknowledged, and strategies in development. For the party that created it, that same log is the digital equivalent of a legal pad covered in case-strategy notes.

In the space of two days this June, courts in two jurisdictions gave the first reasoned answers – and both came down on the side of protection. On June 3, 2026, the Texas Business Court held in *Tate Group Automotive, LLC v. Legacy Automotive Capital, LLC* that a litigant's ChatGPT conversations were protected work product.¹ The next day, June 4, the Supreme Court of New York, Nassau County, quashed non-party subpoenas seeking a litigant's AI prompts, uploads, and outputs in *Assini v. Hayward*, recognizing that AI-assisted litigation preparation can qualify as protected trial-preparation material.² Together, the decisions supply litigators in two major jurisdictions with citable authority for shielding AI-assisted work – and a roadmap for the practices that make protection stick.

Assini v. Hayward: **New York Protects the Pro Se Litigant's AI Workspace**

In *Assini v. Hayward*, the plaintiffs took direct aim at an adversary's AI account. They served a non-party subpoena – and then an amended subpoena – on an AI

provider, seeking materials tied to the accounts of a self-represented defendant: all prompts, inputs, uploaded materials, and corresponding outputs used to draft, revise, or generate filings, motions, sworn statements, or communications transmitted in the case.³

The court granted the motion to quash in its entirety under Civil Practice Law and Rules (CPLR) 2304.⁴ Two features of the ruling deserve particular attention. First, the court held that a self-represented litigant may claim work-product and trial-preparation protection for AI-assisted case preparation. The defendant likened his use of the AI tool to jotting notes on a legal pad or talking through strategy with a colleague, and the court – adopting the reasoning of *Morgan v. V2X, Inc.*⁵ – agreed: confidential, strategy-focused preparation is exactly what the trial-preparation doctrine exists to shield, whether the litigant’s sounding board is a colleague, a notebook, or a chatbot. Second – and of broader consequence – the court held that using a commercial AI tool does not, by itself, waive the protection. The plaintiffs’ implicit theory was that typing case strategy into a third-party service forfeits confidentiality. The court declined to adopt it.

At the same time, the court made clear that AI-assisted litigation is not a supervision-free zone. It underscored its broad authority over disclosure, directed adherence to the New York court system’s new rules on AI use embodied in 22 NYCRR Part 161, and coupled the direction with a sanctions warning.⁶ Protection for AI work product, in other words, coexists with accountability for AI misuse – hallucinated citations and unverified filings remain sanctionable regardless of how the underlying chats are classified.

Tate Group: **Texas Extends Work Product to a Represented Party’s Chats**

The Texas decision reaches further in one important respect. In *Tate Group Automotive, LLC v. Legacy Automotive Capital, LLC*, Judge Grant Dorfman held that ChatGPT conversations conducted by a party principal – not a lawyer, and not a pro se litigant acting as his own counsel – were protected attorney work product under the Texas rule, and that using ChatGPT did not waive the protection.⁷

The defendants pressed *United States v. Heppner*, a federal decision rejecting similar privilege claims, but the court found it inapposite: the question was governed by the

Texas work-product rule, not the federal standard *Heppner* construed.⁸ Texas Rule of Civil Procedure 192.5 defines work product to include material prepared or mental impressions developed in anticipation of litigation by or for a party or a party's representatives – language broad enough, the court concluded, to reach a party's own AI-assisted preparation undertaken in connection with the litigation.⁹

The ruling came with a significant qualification. While the chat logs themselves were shielded, the court ordered the plaintiff to disclose all discovery materials and work products that had been shared with ChatGPT – including materials produced under the protective order. The lesson is double-edged: the strategic content of AI sessions may be protected, but feeding an opponent's confidential documents into a third-party AI tool has consequences, and courts will use their supervisory powers to police what goes into the machine even as they protect what comes out.

The Emerging Doctrine – and Its Limits

Read together, *Assini* and *Tate Group* sketch the outline of an emerging doctrine. AI prompts and outputs created because of litigation – to draft filings, develop strategy, and marshal facts – are trial-preparation materials in substance, and courts applying state work-product rules in New York and Texas have now said so. The medium does not control; the purpose and content do. And in both cases, courts rejected the argument that involving a commercial AI provider automatically destroys confidentiality – a holding with obvious analogues in the treatment of other third-party litigation vendors, from copy services to e-discovery platforms.

But the protection is neither absolute nor uniform, and several fault lines are already visible. Federal courts may take a different path: *Heppner* shows that the federal work-product standard can be applied less generously, and until courts of appeals weigh in, the state-federal divergence will matter enormously to forum strategy. Waiver questions remain fact-intensive: a litigant who uses an AI tool's default settings – under which prompts may be retained, reviewed, or used for model training – presents a harder confidentiality case than one using an enterprise deployment with contractual confidentiality protections, and future opinions will surely probe those details. Attorney-client privilege, as distinct from work product, raises separate problems: a conversation with a chatbot is not a communication with counsel, and litigants should not assume that rules designed to protect

confidential communications with a lawyer will automatically protect exchanges with an AI tool. And where AI outputs are used for purposes beyond case preparation – business decision-making, regulatory compliance, communications later placed at issue – the protection analysis changes entirely.

The protection also stops where a different discovery regime begins: expert practice. In *Conservation Law Foundation, Inc. v. Shell Oil Co.*, a Connecticut federal magistrate judge ordered a party to produce the generative-AI prompts its testifying expert used to cull the document universe underlying her report, reasoning that the prompts were part of the expert’s discoverable methodology under Rule 26 rather than protected notes, drafts, or communications.¹⁰ The order – which, as of this writing, remains stayed pending the district court’s resolution of the plaintiff’s objection – reconciles readily with *Assini* and *Tate Group*: work-product doctrine shields a party’s confidential litigation preparation, but a testifying expert’s methods are affirmatively discoverable so that the opposing party can test the reliability of the resulting opinion. The same idea appears in *Tate Group* itself: the chats were protected, but the court still ordered disclosure of the documents that had been fed into the tool. The dividing line is simple. AI used privately to think through a case is protected; AI used to produce the evidence or opinions a party puts before the court is not. Counsel retaining experts should assume the expert’s AI use will be discoverable and preserve the prompts and outputs from the start.

There is also the “substantial need” back door. Work-product protection, unlike privilege, can be overcome on a showing of substantial need and undue hardship, and ordinary (non-core) work product receives less protection than mental impressions and strategy. Opponents will argue that AI outputs are mere factual compilations subject to the lower tier. Expect the fights to move from “is it protected at all” to “which tier, and is the showing met.”

Practical Guidance for Litigants and Counsel

The decisions reward litigants who treat AI use with the same discipline they apply to other privileged workstreams, and they suggest several concrete practices.

Segregate litigation AI use. Prompts and outputs earn protection because they are created for litigation. Mixing case-strategy sessions with general business queries in

a single account or thread invites line-drawing disputes and partial disclosure. Dedicated matters, dedicated accounts, and clear labeling strengthen the “because of litigation” showing.

Mind the confidentiality settings. Enterprise AI deployments with no-training commitments, retention controls, and confidentiality terms materially improve the waiver analysis. Consumer-grade tools with default data-sharing settings are the weakest footing on which to claim confidential work product – and, as *Tate Group* shows, sharing an opponent’s protected documents with any third-party tool can trigger disclosure obligations of its own.

Counsel clients – including non-lawyers – early. *Tate Group*’s extension of protection to a party principal’s own chats is encouraging, but it is one trial-court decision. Clients should be instructed at the outset of a matter about which AI tools may be used for case-related work, under what settings, and with what materials. Protective orders should be updated to address AI tools expressly, both to restrict feeding produced documents into them and to anticipate discovery requests aimed at AI usage.

Prepare for both sides of the fight. The same doctrine that shields a client’s AI sessions will shield the adversaries. Litigators should calibrate discovery requests accordingly – targeting AI-related information that falls outside protection, such as what produced documents were uploaded to third-party tools – and should be ready to log and defend their own clients’ AI materials with the specificity privilege logs require.

Looking Ahead

Assini and *Tate Group* are trial-court decisions, and the questions they answer will be relitigated in other states, in federal court, and eventually on appeal. But first decisions frame the field, and these two frame it favorably for parties that use generative AI thoughtfully in litigation. Courts, meanwhile, are pairing protection with supervision – New York’s Part 161 rules and the sanctions warnings in both jurisdictions make plain that shielded chats are not a license for unverified filings. For clients and counsel alike, the moment calls for deliberate AI governance in litigation: the protection is real, but it belongs to those who build the record to support it.

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

¹*Tate Group Automotive, LLC v. Legacy Automotive Capital, LLC*, No. 25-BC11B-0020 (Tex. Bus. Ct., 11th Div. June 3, 2026) (minute entry regarding ChatGPT materials in camera review), available at https://websitedc.s3.amazonaws.com/documents/Tate_Group_v._LEgacy_USA_3_June_2026.pdf

²*Assini v Hayward*, 2026 NY Slip Op 26086 (Sup Ct, Nassau County June 4, 2026), available at https://www.nycourts.gov/reporter/current/3dseries/2026/2026_26086.shtml

³*Assini*, 2026 NY Slip Op 26086, available at https://www.nycourts.gov/reporter/current/3dseries/2026/2026_26086.shtml

⁴CPLR 2304, available at <https://www.nysenate.gov/legislation/laws/CVP/2304>

⁵*Morgan v. V2X, Inc.*, 2026 WL 864223 (D. Colo. Mar. 30, 2026)

⁶22 NYCRR Part 161 (Use of Artificial Intelligence Technology), available at <https://www.nycourts.gov/rules/part-161-use-artificial-intelligence-technology>

⁷*Tate Group*, No. 25-BC11B-0020, minute entry at 2-3, available at https://websitedc.s3.amazonaws.com/documents/Tate_Group_v._LEgacy_USA_3_June_2026.pdf

⁸*United States v. Heppner*, 2026 WL 436479 (S.D.N.Y. Feb. 17, 2026)

⁹Tex. R. Civ. P. 192.5, available at <https://www.txcourts.gov/rules-forms/rules-standards/>

¹⁰*Conservation Law Foundation, Inc. v. Shell Oil Co.*, No. 3:21-cv-00933 (D. Conn. May 18, 2026)

(ECF No. 970) (order stayed as of July 2026 pending district court review of plaintiff's objection); see Arnold D. Lee, [Court Orders Disclosure of Expert Witness's AI Prompts: What Litigators Need to Know](#), Spencer Fane (July 8, 2026).

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