Texas Business Court: Work Product Protection Can Extend to a Litigant's AI "Conversations" — but With Important Limits

June 25, 2026

RELATED PRACTICES Trials, Labor and Employment, Litigation

AI and Protected Work Product

A Texas Business Court recently addressed whether a litigant's interactions with ChatGPT qualify for work product protection under the Texas Rules of Civil Procedure. The ruling, one of the first from a Texas state court on this issue, holds that a litigant's generative AI "conversations" prepared in anticipation of litigation can constitute protected work product, but it imposes meaningful transparency obligations on the party claiming that protection.

What Happened

The defendants in *Tate Group Automotive, LLC v. Legacy Automotive Capital, LLC*, No. 25-BC11B-0020, sought to compel production of "conversations" between the plaintiff's executive, Kris Tate, and ChatGPT. The defendants argued that work product protection does not extend to a nonlawyer's AI chats or, alternatively, that sharing information with ChatGPT waived any protection. On June 3, 2026, Judge Grant Dorfman of the Business Court of Texas, Eleventh District, addressed the parties' arguments in a written order.

After an *in-camera* review, the court held that ChatGPT conversations prepared in anticipation of litigation fall squarely within Texas Rule of Civil Procedure 192.5(a)(1), which protects "material prepared or mental impressions developed in anticipation of litigation or for trial by or for a party." The court emphasized that the Texas rule's "by or for a party" language is broader than the federal standard requiring attorney direction and that sharing materials with an AI tool does not waive protection because "work product protections are typically waived by disclosure to an *adversary*, or in circumstances that substantially increase the likelihood that an *adversary* will obtain the materials." (emphasis added).

The court sided with *Warner v. Gilbarco, Inc.* (E.D. Mich. 2026) and *Morgan v. V2X, Inc.* (D. Colo. 2026), and expressly rejected *United States v. Heppner* (S.D.N.Y. 2026), in which the court denied work product protection after the defendant shared materials with Anthropic's Claude platform on his own initiative, without attorney direction and under an Anthropic privacy policy permitting data use for training and disclosure to third parties. The *Tate* court found that distinction significant, but it also ordered meaningful transparency measures.

Key Takeaways

AI work product can be protected — under the right standard. The Texas rule protects materials prepared "by or for a party," which the court found "plainly appear[s] on [its] face to extend

that protection" to nonlawyers. But in jurisdictions applying a narrower, attorney-involved standard (as noted by the *Heppner* court), materials prepared without counsel's involvement may not qualify.

Using AI does not automatically waive protection. The court held that sharing materials with ChatGPT does not waive work product protection because an AI platform is not an "adversary" and disclosure to it does not substantially increase the risk that an adversary will obtain the materials.

But transparency is the price of admission. Despite sustaining the privilege claim, the court ordered the plaintiff to disclose all discovery materials shared with ChatGPT (by Bates number), including materials produced under the protective order. The court further recommended that the parties negotiate amendments to the protective order to address "whether, how, and to what extent" confidential information may be shared with AI tools.

Not everything qualifies. The court ordered production of certain pages from the *in-camera* submission that did not meet the work product standard. The protection applies only to materials genuinely prepared "in anticipation of litigation," not every interaction with a chatbot.

Existing privilege is at risk. While the *Tate* court did not address this issue directly, *Heppner* held that inputting already-privileged attorney communications into a consumer AI tool can independently waive that privilege, particularly where the platform's privacy policy permits data collection, training use and disclosure to third parties. Enterprise AI tools with contractual no-training and no-sharing provisions may be distinguishable, but no court has yet ruled on that distinction.

Expect jurisdictional variation. The court acknowledged that all cited authorities date from 2026, with one calling this "a question of first impression nationwide." The split between *Tate*/*Morgan*/*Warner* (protection available) and *Heppner* (protection denied) reflects a deeper textual divide between party-involved and attorney-involved work product standards. Companies litigating in multiple forums should not assume uniform outcomes.

What You Should Do Now

If your organization uses generative AI in connection with litigation or legal matters, consider these steps:

Update your protective orders. Propose AI-specific provisions at the outset of any litigation. The *Morgan* court crafted model language barring parties from inputting confidential information into AI platforms unless the provider is contractually prohibited from storing inputs for training, disclosing them to third parties and the provider offers the ability to delete confidential information on request. Use this as a starting point.

Keep litigation AI work under attorney supervision. While Texas does not require attorney direction for work product protection, attorney involvement strengthens the litigation-purpose nexus and helps manage waiver risk, a point underscored by *Heppner*, where the absence of attorney involvement was a key factor in denying protection.

Track what goes into AI tools. The court ordered disclosure of all materials shared with ChatGPT, identified by Bates number. Implement logging systems that capture the identity of the tool, the date and what was provided.

Use enterprise-grade platforms. *Heppner*'s waiver finding turned on the AI provider's consumer privacy policy. Enterprise tools with contractual no-training, no-sharing and no-retention provisions may be factually distinguishable, though this theory remains untested.

Do not feed opposing-party confidential materials into external AI. Until protective orders expressly address the issue, treat this as a default prohibition. The court flagged that it would "address any potential violations of the Protective Order's terms if and when they may be shown to have occurred," a clear warning.

This is new territory. The court itself noted that every authority cited dates from 2026. The *Tate* decision offers real assurance that thoughtful use of AI in litigation will not cost you your work product protection. But it also makes clear that courts expect transparency, disciplined information management and proactive protective order provisions. The companies that get ahead of these requirements now will be best positioned as this area of law continues to develop.