Minerva26

# When a Party Feeds Discovery to Generative AI: What *Tate Group v. Legacy* Means for Work Product and Your Protective Order

**Meet and Confer Podcast** · Kelly Twigger **Business Court of Texas, Eleventh Division** · Judge Grant Dorfman · June 3, 2026 🎧 Listen: [PODCAST LINK] · 📄 Case in Minerva26: https://app.minerva26.com/case_law/70496-tate-grp-auto-llc-v-legacy-auto-cap-llc

*Transcript of the broadcast.*

---

Hi, and welcome to the Case of the Week segment of the Meet and Confer podcast. I'm your host, Kelly Twigger.

If you've been watching or listening to Case of the Week this spring, you already know the question that every one of these AI rulings turns on. And it isn't what did the AI do. It's who was at the keyboard, and why. And this week we get to the first state court to pick up that framework and run with it. This decision is from a Texas business court looking at a stack of ChatGPT conversations a company executive ran during litigation, and deciding what's protected, what has to be produced, and what nobody has answered yet.

This case is called *Tate Group Automotive v. Legacy Automotive Capital*. It's from Judge Grant Dorfman in the Business Court of Texas, and the decision is dated June 3rd, 2026. The Texas Business Court is relatively new. The Texas legislature created it in 2023 and it opened in September 2024 as a specialized trial court for high-dollar commercial disputes. So complex business cases get judges with commercial expertise instead of fighting for attention on a district court docket full of criminal, family, and personal injury matters. Part of the point of the new court was to keep businesses litigating in Texas instead of running to Delaware or New York. And the most interesting thing in this particular decision isn't even the ruling, it's the question that the judge flags on the way out the door.

## Where This Case Fits

Let's put this case in context, because the case doesn't really stand alone. If you've been listening, this is part of a long line — a fifth piece in a body of law that's been building since February 2026. Back in February, we had *Warner v. Gilbarco* out of the Eastern District of Michigan, then *Morgan v. V2X* out of Colorado at the end of March. Both involved parties using AI tools in their own litigation, acting as pro se litigants. And both courts said the same thing under the federal work product rule, Rule 26(b)(3). Work product protection applies, and feeding material into an AI tool is not a waiver, because waiver of work product requires disclosure to an adversary, or in a way that's likely to put the material in an adversary's hands. AI is a tool. It is not disclosure to another party. It's not the other side. That's the rule.

Then we had *Heppner* in the Southern District of New York, also in February — came out right about the same time as *Gilbarco*. A represented party in a criminal matter accused of significant embezzlement, using AI outside of and against his own lawyer's direction. There, the court found no protection. *Heppner* is the case that the defense bar reaches for when they want to argue waiver. And that's what happened here.

So coming into June, the map really looked like this. *Morgan* and *Gilbarco* are the rule: a party using AI in support of the case keeps work product protection as a pro se litigant, or where there's advice of counsel. *Heppner* is the exception — when the use case falls outside of the relationship the protection is built to serve, it doesn't hold. *Tate* now tells us which way a court goes when those two lines of authority are sitting on the desk in front of it.

## What Happened in Tate

Here's what happened in *Tate*. Tate Group is the plaintiff. The principal is Kris Tate. During the litigation, Mr. Tate had a series of conversations with ChatGPT, and Tate Group withheld those conversations from production on the basis of attorney work-product. The defendants pushed back on two grounds. First, that work product can't protect a non-lawyer's chats with an AI tool at all. That's *Heppner*. And second, in the alternative, that even if it could, Mr. Tate waived any protection the moment he put the material into ChatGPT. Also *Heppner*, but contrary to *Morgan* and *Gilbarco*.

The defendants also asked for something more. And I want you to hold on to this because it's where the case really gets interesting. They asked the court to order Tate Group to identify all of the discovery materials that Mr. Tate had shared with ChatGPT. That issue came up in *Morgan*, but never really got addressed. The court took the conversations from Mr. Tate in camera. And this minute entry is the ruling that came back from that in camera review.

On the first question, the court sides with Tate, and it does it by name. It adopts the reasoning in *Warner v. Gilbarco* and *Morgan v. V2X*. Work product protections are waived by disclosure to an

adversary — an adversary, that's what's key — or in circumstances that substantially increase the likelihood an adversary will obtain the material. Using ChatGPT is neither of those. So the use of the AI tool is not a waiver in and of itself. It's not a disclosure to an adversary. And the court expressly disagrees with *Heppner*, which is the authority that the defendants built their waiver argument on. So if you're keeping score on the series, *Morgan* and *Gilbarco* have been adopted, so they're in the win category. *Heppner* has been rejected. This is a similar pattern that we're seeing — the rule holds and the exception stays an exception.

Then the court goes one step further. And this is really the part that matters for everyone practicing in state court. It says that *Warner*, *Morgan*, and *Heppner* were all federal cases decided under Rule 26(b)(3). Texas has its own work product rule, Rule 192.5, and that rule is broader. It protects material prepared and mental impressions developed in anticipation of litigation — quote — by or for a party. By or for a party, not by or for a lawyer. Big difference. So broader protection. The use of AI here in this situation extends to Kris Tate as part of Tate Group, because the Texas rule extends the protection of privilege to a party.

So under Texas law, the answer is even easier. Mr. Tate is a party. His litigation-related ChatGPT conversations are the mental impressions of a party developed in anticipation of litigation. They fit inside the rule on its face. The court did order a specific set of pages produced — the ones that, upon in-camera review, weren't actually work product. But the core holding from this decision is that a party's AI conversations about the case are protected, and using the AI didn't waive it in Texas. That's the key.

## What You Still Have to Disclose

Now, here's the turn. Protecting the content of the conversations is one thing. The court then says that work product protection extends only so far, and it reaches back to *Morgan* again. In *Morgan*, the court made the party disclose the identity of the tool he used. Reasoning by analogy, Judge Dorfman here in *Tate* ordered the Tate Group to disclose to the defendants all of the discovery materials that it had shared with ChatGPT, by Bates number where applicable. So a complete listing of everything that it had input into ChatGPT.

Now, read that carefully because there are really two different things happening here, and litigators need to keep them separate. The substance of what Mr. Tate asked ChatGPT — the analysis, his strategy, his mental impressions — that part stays protected. But the fact of which produced documents he fed into the tool has to be provided in discovery, including, in the court's words, quote, any materials that were produced pursuant to the protective order. And that single clause right there is kind of the whole episode here, because that's the next build-on, and the question that the court leaves for us.

## Is Sharing With AI a "Disclosure"?

This is the question that I think every single litigator needs to be carrying into their next protective order negotiation. Mr. Tate took material that the other side had produced under a protective order, and he put it into ChatGPT. The court has now ordered him to identify exactly which protected documents those were. And then the court says something very deliberate. It says it will — quote — be prepared to address any potential violations of the protective order's terms if and when they may be shown to have occurred — close quote. The court's not ruling that a violation has happened yet. It's putting a flag in the ground and saying, tell me what the documents were, whether they were subject to the protective order, and then we'll decide if there was a violation of the protective order.

So think about what a standard protective order actually says. It says a party shall not disclose confidential information to anyone outside a defined list of permitted recipients — counsel, experts, the courts, support staff. That's generally the language. Don't disclose to anyone not on the list. Now ask the question that this case forces. When you paste confidential material into a generative AI tool, have you disclosed it? And that's a real argument either way, and that's exactly the problem.

On one side, of course you have. The text leaves your control. It travels to a vendor's servers. Depending on the tool and the tier, it may be retained, it may be reviewed, it may be used to train a model. That is a disclosure to a third party who is nowhere on the permitted recipients list. On the other side of the argument is, no, you haven't made a disclosure. AI is a tool. You ran the document through software the same way you'd run it through a review platform or store it on a hard drive. The party is using an instrument to do its work in conjunction with it. Nobody would call saving a confidential PDF to your laptop a disclosure. But then again, we wade into the generative part of the AI tool, where the AI tool is giving you content back. Normally, your technology doesn't do that. It may sort, it may filter, it may organize information for you, but it's not physically giving you input on that information.

The reality is that most protective orders on file right now don't answer the question of what disclosure means in the context of generative AI. They were drafted before any of us were thinking about large language models. And the word disclose is now doing work that drafters never contemplated. So whether your client just breached a protective order by using ChatGPT may come down to how a judge reads a word that was written for a world before generative AI. And that's not a comfortable place to be litigating from.

And the judge here knows it, which is why the last substantive thing that he does is recommend that the parties meet and confer to negotiate amendments to the protective order to make it — quote — unquestionably clear whether and how and to what extent confidential information may be shared with any AI tool or large language model. And then he points them straight back to the discussion in *Morgan* for the language.

And this is really the throughline. If we're kind of drawing a line through all these different cases, this is the one that we've been discussing this entire spring, finally landing in one place. *Morgan* gave us the clauses: no training on your inputs, no onward disclosure, deletion on demand. *Tate* gives us the reason you can't wait for the other side to raise it, because the ambiguity doesn't sit in a neutral place. It sits as exposure on whichever party moved first without a clear rule. And in this case, that party is the one that used the tool.

## Takeaways

All right, let's talk about the takeaways and what this means that you need to work on right now. First, the protection on these facts is real, but it is not a license, because you need to be precise about why this Texas court held what it did — because the reason is the rule in Texas on privilege, and not anything to do with the AI tool. Work product protects a party's own use in three situations. One, where the party is pro se, because then the party is effectively counsel. That's *Morgan* and *Gilbarco*. Two, the party is acting on the instruction of counsel, because then the work is done for the lawyer. Those two come from federal law, Rule 26(b)(3). Three, you're in a jurisdiction whose rule protects work prepared by or for a party. That's Texas Rule 192.5, that we're dealing with here. And that's why *Tate* came out the way it did. Three separate situations. Know them, understand them.

Now, the trap here. Mr. Tate was a represented party using AI on his own, not pro se and not, on this record, at the direction of counsel. Under federal law, that profile looks like *Heppner*, and *Heppner* lost. But *Tate* is protected here for only one reason, and we just talked about it: the Texas rule includes the word party, and it protects work product on the face of the rule. That language is an outlier, and it's not the norm. So if your matter is in federal court, or in a state that tracks the federal standard, do not assume that *Tate* covers you. There's still an analysis to run. I would also caution you against using *Tate* as leverage in a federal matter, because a court may not appreciate you disingenuously providing case law on a rule that is different than the federal standard.

Second, assume you will have to account for what went into the platform. The content of the conversation may be privileged, but the inventory of documents you fed the tool is increasingly not. And we'll talk in a second about what you get out of knowing what the documents are that the other side put in. The inventory of documents can tell you what they find important, what kind of questions were asked, and then you can conduct a deposition about why those documents were important to them. If a produced document touched an AI platform, you want to start building the record now. What was used, on what platform, on what data? Document, document, document — one of our three rules of Case of the Week. You've heard me say it all the time, and it has never been more literal than it is here. You need to keep a list of the documents that you're putting into AI so that you can validate that you're complying with the protective order.

Third takeaway, and this is the strategic one. Go and look at the protective orders in your active matters and ask whether the word disclose covers an AI tool. If you can't answer that with confidence, neither can the judge. And that uncertainty is a liability that you're carrying for free right now. The fix is the same one in *Morgan* that *Tate* is suggesting. Bring AI-specific language to the table for your protective orders. Go back and amend your existing protective orders if necessary. Don't wait to negotiate it from behind after a stack of confidential documents has already gone into a tool.

Now, there's a flip side to all of this, and it's something that I raised when we first talked about the *Morgan* case, and that is that Rule 34 historically has not required us to provide this kind of information in terms of what we do to provide information in production under Rule 26, and whether we use TAR — those kinds of things. So those questions are still totally unanswered. There are cases on both sides of the equation and no definitive rule about what has to be disclosed. And so there's a question here for me about disclosing this AI usage. And if you don't negotiate it into your protocol, and if one side has an enterprise-grade AI tool, so they're meeting all of the obligations of *Morgan* in terms of protecting the data, but there's still no specific language in a protective order — so it's kind of like the debate about an ESI protocol. Do I do it? Do I not do it? You need to be carefully considering that. I'm telling you what the law is saying right now in these various courts. Obviously, if you're in a different court than any one of these decisions, they're not precedential. But as we've seen in every issue in discovery case law on ESI, the district courts across the country are looking to the other district courts across the country that have ruled on these specific issues for guidance. They're looking to try and create consistency in the law on these issues. So precedential or not precedential, there's going to be value in these decisions, and you're going to have to think about exactly what steps you're taking to leverage AI as a tool in discovery and how it's going to interplay with protective orders.

Now, because the pattern across all five of these cases is the same one that we keep coming back to: the courtroom is converging on accountability faster than firms are updating their playbooks. The litigators who are going to be fine are the ones who can answer, for every active matter, who's using AI, in what role, on what platform, and with what data. This is now something that you're probably going to need to add to your custodial interviews. What are the AI sources that those custodians are using? Talk to your experts about it. It's got to be on your list now of what's happening. And that's because it's not a tool to filter and sort data. It's giving you output that then becomes a question as to whether or not that is discoverable. Put in writing the answers to these questions before they get asked. It's just like doing all the documentation that you do when you're in the process of interviewing custodians, identifying data sources, making decisions about what to preserve and what to collect. This is another piece of that puzzle. So put all of this information in. Make sure you know and understand what tools your clients have available to them. If they're using personal ones, you're going to have to tell them this is — it's all going to be potentially discoverable. And we don't want to get in the situation where we have to produce it under a motion to compel, because then we have very little control over how we can make arguments. So get ahead of this. Document, document, document — one of our key themes here. Plan, start early, document it. All three of these takeaways are

strategic decisions, and all three have to be made by the litigators who are going to live with the consequences and by the clients who are going to pay for them.

## Resources and Final Notes

Now, this line of cases is moving very quickly and they're not slowing down. We've seen probably more than 40 or 50 decisions on protective orders related to generative AI, and new protective orders coming in with AI language in them to deal specifically with disclosure to AI tools, in Minerva26. So if you are interested in following this line of cases as it develops, you can leverage the Generative AI tag in Minerva26 to be able to receive updates of those cases as they're coming about. If you want to follow specifically the protective orders with generative AI, you can search for issue tags Generative AI and Protective Order to be able to receive those decisions delivered directly to your inbox. So when you sit down to draft that protective order amendment, the whole body of authority is already pulled together for you instead of buried in some search. And that's the difference between really reacting to this shift and getting ahead of it. So spend a couple of hours looking at those cases and the new protective order language, update the forms that you're using for every matter, look at your active litigations and determine where do you have potential issues. If you've got a case where discovery's closed and it's five years old and it may be heading to trial, or you're trying to resolve it, that may not be the case where you need to amend your protective order. But if you've got an active discovery matter going on right now with the potential for folks using AI, that's where you're going to want to revisit this question.

That's it for our Case of the Week this week. One other announcement to be aware of. This week, the University of Florida Journal of Law, Technology and Policy released its special eDiscovery edition that was developed in partnership with the University of Florida eDiscovery Conference. I'll include the link to all of the articles in that journal, and you'll want to check them out. The content is exceptional. There are two pieces in particular that I think you'll want to check out — from United States District Judge Xavier Rodriguez, and another one from United States Magistrate Judge William Matthewman. They're completely free at the link to read. And the authors in this edition are some of the most forward thinkers that we have in this profession. And they are practical. Read them. The article from Robert Keeling on looking at the use of generative AI in document review is incredibly compelling. It's going to give you a lot of information and a leg up on trying to figure out these issues on your own.

I'm Kelly Twigger. Thanks so much for joining me. If this episode was helpful, please share it with a colleague. These issues are multiplying faster than the mosquitoes in my backyard, and every litigator and in-house counsel needs to be aware of how quickly the law is developing. The more people we reach, the better everyone does discovery. Wishing everyone a fantastic Fourth of July celebration, and here's hoping that our democracy survives for another 250 years. I'll be on vacation the week after the Fourth, so I'll be back later in July with the next update on what's happening in the world of discovery. Have a great week.