LAW360

Portfolio Media. Inc. | 230 Park Avenue, 7th Floor | New York, NY 10169 | www.law360.com
Phone: +1 646 783 7100 | Fax: +1 646 783 7161 | customerservice@law360.com

Tax Teams Get No Bright-Line Rule From AI Privilege Cases

By **Drew Cummings and Jennifer Breen** (May 22, 2026, 3:49 PM EDT)


Drew Cummings


Jennifer Breen

In Morgan v. V2X Inc., the U.S. District Court for the District of Colorado offered the newest signal on how courts evaluate privilege and work-product issues when artificial intelligence is used, on March 30 holding that the use of AI, by itself, would not defeat the work-product protection. The court held that a pro se litigant's AI-assisted materials fell within Rule 26(b)(3), while also imposing practical guardrails around disclosure of confidential information to an AI tool.

February decisions from the U.S. District Court for the Eastern District of Michigan in Warner v. Gilbarco Inc.[2] and the U.S. District Court for the Southern District of New York's Feb. 17 decision in U.S. v. Heppner remain important as the contours around AI privilege start to take shape.[3]

The practical question now is what these signals mean for companies, including their tax departments, as they deploy AI in day-to-day operations. Until greater clarity emerges, tax practitioners and taxpayers alike should approach AI use thoughtfully and evaluate how those tools may affect the protection of sensitive materials.

To understand the significance of Morgan, however, it is useful to begin with Heppner, where the court addressed the threshold question whether communications with a publicly available AI platform could have privilege attach in the first instance, and whether such communications could then retain either privilege or work-product protection.

U.S. v. Heppner

In Heppner, the Southern District of New York addressed whether communications with a publicly available AI platform can be protected by the attorney-client privilege or the work product doctrine, and held that they could not. The court's analysis focused on core privilege requirements, concluding that the defendant's use of an external AI tool undermined the elements necessary to sustain protection, particularly with respect to confidentiality and in the absence of an attorney relationship.

Why were the AI communications not protected by attorney-client privilege?

The court concluded that "at least two, if not all three," of the following elements were lacking.

First, the communications were not between a client and an attorney. The prompts and AI-generated responses consisted of interactions with a publicly available AI platform, which the court determined could not be treated as a lawyer. The fact that Heppner later shared the AI-generated materials with his defense counsel did not transform his AI communications into privileged ones.

Second, the court found no reasonable expectation of confidentiality. The platform's privacy policy allowed user inputs to be used for training and disclosed to third parties, including government agencies, undermining any claim that the communications were intended to remain confidential.

Third, the court questioned whether the communications were made for the purpose of obtaining legal advice. The court noted that Heppner used the AI tool on his own initiative, rather than at the direction of counsel.

In the alternative, the court held that even if privilege had initially attached, it was waived when Heppner disclosed the information to the AI platform.

Why were the documents containing the AI communications not protected by the work product doctrine?

The court held that the documents containing the AI communications were not protected because they were created by Heppner alone and not at the direction of counsel, despite the fact that Heppner's counsel had provided him with the information that he input into the AI platform in the first instance.[4]

Warner v. Gilbarco

In Warner, the Eastern District of Michigan reached a different result, denying a request for production of documents and information reflecting a party's use of third-party AI tools. The court held that such materials were protected work product, emphasizing that the doctrine extends to materials prepared by or for a party in anticipation of litigation, even where open-source AI tools are used.

Critically, the court rejected the argument that use of an AI platform effected a waiver, concluding that there was no showing that the materials were disclosed "to an adversary or in a way likely to get in an adversary's hands."

Morgan v. V2X

In Morgan, the District of Colorado similarly declined to treat AI use as inherently incompatible with work-product protection. There, the court held that a pro se litigant's use of AI could fall within Rule 26(b)(3), reaffirming that work-product protection applies to materials prepared by a party, even in the absence of counsel.

The court distinguished Heppner based upon both procedural context (criminal vs. civil litigation) and the absence of any gap between the litigant and counsel. At the same time, the court imposed guardrails, requiring disclosure of the AI tool used and restricting the submission of confidential information to systems lacking contractual safeguards.

Taken together, these decisions do not establish a bright-line rule but instead point toward a fact-intensive framework, focused on confidentiality and disclosure risk.

Implications for Tax Departments

In the tax context, privilege protections are foundational. The attorney-client privilege, the Internal Revenue Code Section 7525 tax-practitioner privilege and the work product doctrine each protect sensitive materials from disclosure.

Courts have long applied privilege doctrines to evolving technologies, including email, text messages, cloud storage and electronic document management systems, which all led to early uncertainty before more-settled principles emerged.[5] Generative AI presents the next iteration of that challenge. Against this backdrop, tax departments should carefully evaluate how their use of AI tools may affect privilege protections.[6]

Waiver Risk of Publicly Available AI Platforms

At a minimum, the Heppner decision underscores the risk that uploading privileged material to a publicly available AI platform may constitute a waiver of the attorney-client privilege. The court's alternative holding — that even if privilege attached it was waived by disclosure to the AI platform — should alone give tax departments pause.

Consider a common scenario: An employee uploads a privileged legal opinion or memorandum that is protected by the attorney-client privilege or Section 7525 and asks the AI tool to generate counterarguments or risk assessments. Under the available case law, such an act may constitute waiver if a publicly available AI model is utilized.

To be fair, Heppner does not stand for the proposition that AI use categorically results in a waiver, especially with respect to the work product doctrine. In Warner and Morgan, the courts rejected a waiver argument and upheld work-product protection where there was no indication that the use of an AI platform would result in disclosure to an adversary.

Taken together, these decisions suggest that the relevant inquiry is not simply whether AI is used, but whether its use introduces a third party in a manner that is inconsistent with maintaining confidentiality or creates a meaningful risk of disclosure. There is authority suggesting that disclosure to third-party technology providers does not automatically result in waiver where confidentiality is preserved through contractual protections.[7]

Low or no-cost AI tools, however, may present a different risk profile at the moment, particularly where privacy policies allow for data use, model training or discretionary disclosure.[8] On the other hand, enterprise AI systems — nonpublic tools with restricted access and strong contractual protections — may present lower risk.[9]

Even so, the law remains unsettled, and the use of any AI system should be evaluated based upon how the tool operates, how data is handled and whether confidentiality is preserved.

Practical Guardrails for Tax Departments

Privilege determinations are inherently fact specific. In light of these decisions, tax departments should consider implementing the following guardrails.

Restrict use of public AI tools for sensitive matters. Avoid uploading privileged opinions, memoranda or

litigation analyses into publicly available AI tools, and consider implementing technical controls — e.g., blocking or limiting access — to prevent inadvertent uploads of privileged or confidential information into high-risk platforms.

Minimize and anonymize inputs. Avoid including identifying or sensitive factual details unless necessary. Use redacted or hypothetical inputs where possible.

To the extent possible, have analyses generated under the supervision of in-house or outside counsel or a federally authorized tax practitioner, rather than by nonlawyers or by tax practitioners who are not authorized or that act independently.

Limit AI-assisted legal analyses in sensitive matters.

Confirm the basis for work-product protection. Ensure there is a reasonable anticipation of litigation before generating AI-assisted analyses.[10]

Review vendor terms and system protections. Confirm that enterprise AI systems provide sufficient confidentiality and restrict data use.

Implement clear internal policies governing AI use. Establish written protocols governing AI use and ensure personnel understand privilege risks. Update internal document retention policies to include AI platforms.

Maintain documentation of AI use and safeguards. Be prepared to demonstrate how the tool operates, what protections are in place and why the use did not create a meaningful risk of disclosure, as this can be critical in defending privilege claims.

Coordinate with information technology and legal on system architecture. Evaluate whether AI tools store, retain or transmit data externally, and align usage with systems that support a defensible expectation of privacy.

Looking Ahead

As courts begin to address how traditional privilege doctrines apply to generative AI, uncertainty will persist. Decisions such as Heppner, Warner and Morgan illustrate that existing principles can yield different outcomes depending on how AI tools are used and how information is handled. Because the privilege protections are all inherently fact specific, taxpayers and tax practitioners should evaluate AI use on a case-by-case basis, with particular attention to confidentiality, disclosure risk and system design.

Tax departments should adopt disciplined AI practices, including clear policies, training and defined limits on the use of public tools. Until further clarity emerges, they should assume that the use of publicly available AI platforms may jeopardize privilege and structure their workflows accordingly.

Drew A. Cummings is a partner at Morgan Lewis & Bockius LLP.

Jennifer Breen is a partner at the firm. She previously served as an attorney with the IRS Office of Chief Counsel's National Office

The opinions expressed are those of the author(s) and do not necessarily reflect the views of their employer, its clients, or Portfolio Media Inc., or any of its or their respective affiliates. This article is for general information purposes and is not intended to be and should not be taken as legal advice.

[1] Morgan v. V2X Inc., 2026 WL 864223 (D. Colo. March 2026).

[2] Warner v. Gilbarco Inc., No. 2:24-cv-12333, at 11 (E.D. Mich Feb. 10, 2026).

[3] United States v. Bradley Heppner, 25 CR 503 (JSR), Bench Memorandum dated February 17, 2026 (Document 27).

[4] But see Tremblay v. OpenAI Inc., 2024 WL 3748003 (N.D. Cal. Aug. 8, 2024) (finding that work product attached to ChatGPT prompts and outputs crafted by counsel that contained counsel's mental impressions but that certain prompts and outputs affirmatively used in the case had been waived).

[5] See, e.g., In re: Asia Global Crossing Ltd., 322 B.R. 247 (S.D.N.Y. 2005) (evaluating email confidentiality); see also ABA Formal Opinion 99-413 (questioning unencrypted email confidentiality). See also Quon v. Arch Wireless Operating Co. Inc., 529 F.3d 892, 905 (9th Cir. 2008), rev'd 560 U.S. 746 (text message privacy); Harleysville Ins. Co. v. Holding Funeral Home Inc., 2017 WL 4368617, at *7 (W.D. Va. 2017) (cloud storage safeguards); Diamond Resorts U.S. Collection Development LLC v. US Consumer Attorneys PA, 2021 WL 4482913 (S.D. Fla. 2021) (no waiver via upload to document management system).

[6] When not subject to privilege protection, AI prompts and outputs are already being used affirmatively at trial against parties. For example, in Fortis Advisors LLC v. Krafton Inc., 2026 WL 730977 (Del. Ch. Mar. 16, 2026), Krafton's CEO's ChatGPT prompts and responses were entered into evidence.

[7] See, e.g., Harleysville Ins. Co., 2017 WL 4368617, at *7; State Bar of Arizona, Opinion No. 05-04; State Bar California Standing Committee on Professional Responsibility and Conduct, Formal Opinion No. 2012-184.

[8] Besides privilege waiver concerns, there are also confidentiality concerns with respect to uploading highly sensitive business information into publicly available AI models — what if those materials are somehow made public or used by third-parties.

[9] See New York City Bar Formal Opinion 2025-6 (noting the impact of AI's use on confidentiality).

[10] Courts have not identified a precise starting point for satisfaction of the "anticipation of litigation" standard in the tax context. Rather, courts have looked to whether litigation was reasonably contemplated at the time the challenged materials were created. Litigation may reasonably be anticipated not merely prior to audit, but even prior to the event (i.e., the transaction) that the taxpayer believed ultimately would give rise to litigation. See United States v. Adlman, 134 F.3d 1194 (2d Cir. 1998), United States v. Chevron Texaco Corp., 241 F. Supp. 2d 1065 (N.D. Cal. 2002.