Insights: Alerts

Recent Decisions Spark Questions on Generative AI, Privilege, and Privacy Expectations

April 23, 2026

Written by **Joel D. Bush**

Three recent federal court decisions address whether materials created using public generative AI platforms are protected by the attorney-client privilege or work product doctrine. The rulings also raise important questions about privacy expectations when using AI tools.

What the Courts Have Said

***United States v. Heppner* (S.D.N.Y. Feb. 17, 2026).** *Heppner* involved a former CEO, Bradley Heppner, who was indicted on criminal fraud charges. Heppner used Anthropic's Claude (a publicly accessible platform) to prepare 31 documents related to his legal defense and subsequently sent those documents to his legal counsel. At issue was whether these AI-generated documents were protected by the attorney-client privilege or the work product doctrine.

The court said no:

- Attorney-client privilege did **not** apply because the documents were communications between the defendant and the AI tool, not between the defendant and his counsel.
- In response to a query, Anthropic's Claude responded: “I'm not a lawyer and can't provide formal legal advice or recommendations” and went on to recommend that a user “should consult with a qualified attorney who can properly assess your specific circumstances.”
- The court found that, because the AI Documents “would not be privileged if they remained in [Heppner's] hands,” they did not “acquire protection merely because they were transferred” to counsel.
- According to the court, “non-privileged communications are not somehow alchemically changed into privileged ones upon being shared with counsel.”
- The communications with the AI platform were **not confidential**; Anthropic's user policy permits the company to collect and potentially disclose user data, including inputs and outputs, to third parties such as regulators.
- Work product protection also did **not** apply because the documents “were prepared by the defendant on his own volition,” and thus reflected the client's and the AI's theories -- not counsel's legal strategy.

- The judge noted that his decision might have been different if counsel had directed the use of the AI tool, but that was not the case before the court.

***Warner v. Gilbarco* (E.D. Mich. Feb. 10, 2026).** *Warner* involved a pro se plaintiff who used a public gen AI platform in connection with her lawsuit.

The court held that:

- Materials prepared by a pro se litigant with a gen AI platform, and in anticipation of litigation, **were** protected by the work product doctrine, because a pro se litigant has the right to assert work product protection over such material.
- Disclosures of mental impressions to a public gen AI platform did not constitute waiver because waiver of work product materials requires disclosure to an *adversary* (or at least in a manner likely to end up in an adversary's hands).
- In short, sharing mental impressions with an AI tool is fundamentally different than handing them to opposing counsel.
- According to the court, "ChatGPT (and other generative AI platforms) are *tools, not persons,* even if they have administrators somewhere in the background." (emphasis in original).
- The court observed that, if the use of gen AI is found to waive work product protection, such decision would "nullify work product protection in nearly every modern drafting environment, a result no court has endorsed."

While the result in *Warner* appears inconsistent with the work product conclusion in *Heppner,* the cases involved different circumstances. *Heppner* – unlike *Warner* – involved a defendant who *was represented* by counsel and who was not acting pro se.

***Morgan v. V2X, Inc.* (D. of Colo. Mar. 30, 2026).** In *Morgan,* a *pro se* employment discrimination plaintiff used a gen AI platform in connection with his lawsuit, and the defendant challenged the plaintiff's assertion of work product protections over the outputs from his AI use.

The court held that:

- Like in *Warner,* a pro se plaintiff can assert work product protection in connection with materials generated by an AI platform.
- "In the context of a *pro se* litigant's use of AI to assist with their litigation preparation, the use of AI closely resembles the kind of confidential, strategy-laden iterative work product that Rule 26(b)(3) was designed to protect."
- The court distinguished *Heppner* as involving "a gap between the party and the attorney because the defendant [in *Heppner*] acted entirely apart from his lawyer." "No such gap exists in the *pro se* context" because "[a] *pro se* litigant is simultaneously the party and the advocate."

What These Decisions Mean

On one hand, none of these decisions are remarkable, because they apply well-settled principles. If there is no confidentiality associated with inputs into a public generative AI platform, attorney-client privilege cannot apply. It is long-established that confidentiality is one requirement for attorney-client privileged communications. *United States v. Mejira,* 655 F.3d 126, 132 (2d Cir. 2011). Work product protection extends to pro se litigants. And, while the attorney-client privilege is waived by any disclosure to a third party, work product protection is waived only by disclosure to a litigation adversary. "Disclosure" to a generative AI platform is not disclosure to an adversary. In short, these rulings largely follow from established doctrine applied to new technology.

On the other hand, these cases raise questions about whether an expectation of privacy does – or should – apply to use of public generative AI platforms. Judge Rakoff in *Heppner* concluded that there is no such privacy expectation with use of these public platforms: "AI users do not have substantial privacy interests in their 'conversations with … publicly accessible AI platform[s]'" because "users voluntarily disclose[ ]' [content] to the platform" and "the platform 'retains [the content] in the normal course of its business.'" As a consequence, according to Judge Rakoff, there could be "no 'reasonable expectation of confidentiality in … communications' with Claude."

But Magistrate Judge Braswell in *Morgan* sees things differently:

> "It is true that AI systems like ChatGPT, Claude, Gemini, and others … collect user data for training and other purposes. But … that does not eliminate all expectations of privacy or automatically waive protections."

Magistrate Judge Braswell places gen AI in the broader context of electronic communications and the fact that "our phones, computers, in-home smart devices, and other electronics, collect information about us to offer more bespoke services." She asks whether this reality means "that anyone with a Gmail account has forfeited all rights to confidentiality and privacy?" She observes that, "given how AI tools function, it is entirely reasonable for a person to expect some privacy and confidentiality when interacting with these tools, even though they understand a third party is behind the tool collecting and storing their information."

While these observations are *dicta,* they prefigure the forthcoming debate about whether the routing of sensitive information through a third-party technology system, such as a Google account or a generative AI platform, destroys privacy protections.

What Critics Are Saying

The *Heppner* decision has sparked significant debate, with critics arguing that the ruling overstates the risks associated with generative AI and may mischaracterize the practical application of privilege in the digital age. Notably, an op-ed piece in the *Wall Street Journal*, published on April 6, 2026, argued that "if this reasoning stands, the consequences will reach far beyond artificial intelligence."

Critics contend:

- Judge Rakoff wrongly treated the AI platform "like a person" and having engaged in "communications" with the defendant. "But AI isn't a person … [i]t can't be deposed, call the police or betray a confidence."
- Third-party disclosure risks "don't exist when the 'third party' is a statistical model running on a server."

- Inputting information into an AI platform “is no different from typing into a cloud-based software, such as Google Docs.” The question is not “whether Google Docs creates privilege,” the issue is “whether Google Docs destroys” privilege – and no one thinks that is the case.

The *Heppner* ruling is thus generating concern about the loss of privilege when using AI, potentially discouraging legitimate and efficient legal workflows. Some legal observers believe courts should adapt privilege doctrines to reflect the evolving realities of technology-assisted legal practices.

What You Should Be Doing Now

Given the court's clear stance in *Heppner*, legal professionals and clients should take proactive steps to protect privilege when using AI tools:

- **Avoid Using Public AI Platforms for Privileged Work**: Do not draft sensitive documents or legal analyses on public AI tools where user data may be accessed or disclosed by third parties.
- **Engage Counsel Early**: Ensure that any materials intended to be privileged are prepared at the direction of counsel and as part of a legal strategy, not independently by clients.
- **Review and Update Internal Policies**: Educate clients and teams on privilege risks and best practices regarding AI use. Implement internal guidelines restricting the use of generative AI for privileged or confidential work unless protected by adequate safeguards.

By adopting these precautions, legal departments and firms can minimize the risk of waiving privilege and ensure compliance with current judicial expectations.

We will continue to monitor the evolving case law and guidance from courts regarding privilege and AI use. Please contact us if you have questions about how these decisions may affect your organization.

Related People

Joel D. Bush

t 404.815.6074

jbush@ktslaw.com