FINTECH LAW REPORT
E-Banking, Payments & Commerce in the Mobile World

July/August 2026 ▪ Volume 29 ▪ Issue 4

WHEN A NON-LAWYER OFFICER OR DIRECTOR ASKS AN AI PLATFORM FOR LEGAL ADVICE ON A COMPANY MATTER, ARE THE CONVERSATIONS PROTECTED FROM DISCOVERY IN LITIGATION?—*FORTIS v. KRAFTON*

By Gail Weinstein, Philip Richter, Steven Epstein, Steven Steinman, Randi Lally and Colum Weiden

Gail Weinstein is a senior counsel in the M&A and Private Equity Practice at Fried, Frank, Harris, Shriver & Jacobson LLP. Philip Richter is a partner and co-head of Fried Frank's M&A and Private Equity Practice. Steven Epstein is the managing partner of Fried Frank and co-head of its M&A and Private Equity Practice. Steven Steinman is a partner and co-head of Fried Frank's M&A and Private Equity Practice. Randi Lally is a partner and co-head of Fried Frank's M&A and Private Equity Practice. Colum Weiden is a partner in Fried Frank's M&A and Private Equity Practice. All are based in the New York office. Contact: gail.weinstein@friedfrank.com or philip.richter@friedfrank.com or steven.epstein@friedfrank.com or steven.steinman@friedfrank.com or randi.lally@friedfrank.com or colum.weiden@friedfrank.com.

When a non-lawyer officer or director asks an AI platform for legal advice on a company matter, are these conversations protected from discovery in litigation? One might expect the response to be uncomplicated—say, that such conversations would not be protected from discovery, under either the attorney-client privilege or the attorney work product doctrine, because AI is not *an attorney*. But courts are just beginning to grapple with this question, and the answers have been varied:

- In *U.S. v. Heppner,*[1] a federal district court in New York held that a criminal defendant's exchanges with a consumer version of Claude, which were not directed by his lawyer, were discoverable.
- And, in *Fortis Advisors v. Krafton*,[2] the Delaware Court of Chancery considered as evidence a CEO's ChatGPT exchanges that provided a legal strategy for the company to avoid having to pay an earnout obligation.
- However, in *Warner v. Gilbarco Inc.,*[3] a federal district court in Michigan held that a *pro se* litigant's use of AI was work prod-

IN THIS ISSUE:

When a Non-Lawyer Officer or Director Asks an AI Platform for Legal Advice on a Company Matter, Are the Conversations Protected from Discovery in Litigation?—*Fortis v. Krafton* 1

Bank-Like Customer Identification Program Rules Proposed for Stablecoin Issuers 6

The Next Digital Trade Framework: What an EU-CPTPP Agreement Signals for Data, Cyber and AI Governance 9

CFTC Wants to Know: What Are the Regulatory Barriers for Fintech Firms? 16

FCA Finalizes Core Rules for the UK Cryptoasset Regime 18

AI and Monetary Policy 22

FinTech Law Report: May-June 2026 Regulation and Litigation Update 25

Thomson Reuters™

43373673

uct, and so was protected from discovery, because it was used in anticipation of litigation and in a manner not likely to get into an adversary's hands.

- Also, in *Morgan v. V2X, Inc.*,[4] a federal district court in Colorado held that a *pro se* litigant's use of AI was work product—but also held that the litigant had to disclose which AI tool he was using and that the protective order could permit AI use only where the terms of use protected confidentiality.

- And, in *Tate Group Automotive v. Legacy Automotive Capital*,[5] the Texas Business Court held that litigation materials the plaintiff had uploaded to ChatGPT were protected from discovery as work product—but also held that the plaintiff had to identify (although not disclose the contents of) every document that he had uploaded to the AI platform.

The law on this issue is unsettled and evolving. What has become clear so far, though, is that: (i) when an AI platform is used by a non-lawyer to obtain legal advice about a company matter, there is a significant risk that the queries and output may be discoverable in litigation; (ii) such evidence can be very damaging to a litigant by indicating its motivations or mental state, thought process, risk analysis, and so on; and (iii) the judicial outcome will depend on the particular court's perspective on AI and the specific facts and circumstances of the case.

The key divide in judicial perspective on AI has been whether it should be viewed merely as a litigation "tool" (analogous to, say, Westlaw) or, instead, as being the equivalent of a "person" (an attorney) providing legal advice. The facts and circumstances that the courts have viewed as important have included whether the exchanges with AI were directed by legal counsel; what the particular AI platform's terms of use were, especially relating to confidentiality; whether AI was used in a manner that the exchanges could reach the hands of the litigation adversary; and the specific wording of the federal or state rules of civil procedure that were applicable.

Below, we summarize the courts' decisions and note relevant considerations for crafting corporate policies and protocols.

SUMMARY OF DECISIONS

***Heppner*—AI interactions were not protected.**

In *Heppner*, a New York federal court held that the advice that a criminal defendant received from Claude (a generative AI platform) on legal defense strategy was not protected by the attorney-client privilege or the work product doctrine. The defendant had shared Claude's output with his defense attorney, and the attorney had been "influenced" by (but had not followed) Claude's advice in crafting his legal strategy. The attorney con-

FinTech Law Report
West LegalEdcenter
610 Opperman Drive
Eagan, MN 55123

©2026 Thomson Reuters

For authorization to photocopy, please contact the **Copyright Clearance Center** at 222 Rosewood Drive, Danvers, MA 01923, USA, http://www.copyright.com, Toll-Free US +1.855.239.3415; International +1.978.646.2600 or **Thomson Reuters Copyright Services** at 2900 Ames Crossing Rd, Suite 100, Eagan, MN 55121, USA or copyright.west@thomsonreuters.com. Please outline the specific material involved, the number of copies you wish to distribute and the purpose or format of the use.

This publication was created to provide you with accurate and authoritative information concerning the subject matter covered; however, this publication was not necessarily prepared by persons licensed to practice law in a particular jurisdiction. The publisher is not engaged in rendering legal or other professional advice and this publication is not a substitute for the advice of an attorney. If you require legal or other expert advice, you should seek the services of a competent attorney or other professional.

Copyright is not claimed as to any part of the original work prepared by a United States Government officer or employee as part of the person's official duties.

One Year Subscription ● 6 Issues ● $ 1020.00

ceded that he had not directed the defendant to consult with AI but stressed that the defendant had consulted with AI for the "express purpose of talking to counsel." Judge Jed S. Rakoff stressed that the defendant had no reasonable expectation of confidentiality given the AI platform's terms of use, which included users' consent to their queries and Claude's responses being used by Anthropic (Claude's creator) to "train" Claude and to make disclosures to third parties such as government regulators or in connection with disputes or litigation.

With respect to the attorney-client privilege, the court reasoned that: (i) Claude was not a person and thus "not an attorney"; (ii) the defendant had no reasonable expectation of confidentiality in his communications with Claude given the platform's terms of use; and (iii) the defendant could not have intended to obtain "legal advice" from Claude, as Claude's responses expressly disclaimed that Claude was a lawyer or could provide legal advice. Notably, the court indicated that the result may have been different if the defense attorney had directed the defendant's interactions with Claude—in that case, Claude "might arguably be said to have functioned in a manner akin to a highly-trained professional who may act as a lawyer's agent within the protection of the attorney-client privilege." With respect to the work product doctrine, the court reasoned that, although the exchanges with Claude may have been "in anticipation of litigation," they were not "prepared by or at the behest of counsel"; and, further, they were not reflective of the legal strategy the counsel actually used.

***Krafton*—AI interactions were critical evidence.**

In *Krafton*, a CEO sought, and largely followed, legal advice from ChatGPT as to how to avoid making an earnout payment to the key employees of an acquired company. ChatGPT advised that, legally, it would be difficult to avoid making the earnout payment, but that the CEO could create pressure points by terminating the key employees for cause and, if that was unsuccessful, by effecting a "take over" of control of the acquired company (although the acquisition agreement provided for control by the key employees during the earnout period). ChatGPT suggested formation of an internal task force, provided a name for the project, and provided detailed talking points, letters, and other communications to be used in each step of the process. In the litigation that ensued, the terminated key employees obtained discovery of the CEO's exchanges with ChatGPT and introduced them in evidence. The issue of privilege was not raised, presumably because the CEO had shared the exchanges with other non-lawyers at the company. The Delaware Court of Chancery, matching ChatGPT's advice with the CEO's subsequent actions, found that the key employees were terminated for pretextual reasons and that the takeover of control violated the acquisition agreement. Vice Chancellor Lori W. Will ordered that the CEO be restored as CEO, with operational control, and that the earnout period be extended by 258 days. Money damages are still to be determined.

***Warner*—AI interactions were protected**.

In *Warner*, a Michigan federal court held that the *pro se* plaintiff's interactions with ChatGPT to obtain legal advice were protected under the work-product doctrine. The plaintiff claimed racial discrimination by her employer. During discovery, the defendants moved to compel production of "all documents and information concerning [the plaintiff's] use of third-party AI tools in connection with [the] lawsuit." The plaintiff asserted work product protection. The defendants argued that the plaintiff had waived that protection when she input litigation materials into a public version of an AI platform.

The court noted that the Federal Rules of Civil Procedure (Section 26(b)(3)(A)) shield "documents and tangible things that are prepared in anticipation of litigation or for trial by another party or its representative." A plain reading of that text, the court stated, indicates that the work product doctrine, as articulated in the Federal Rules, protects materials prepared "by [a] party" to the litigation (not just by or at the direction of an attorney). The court stated that the plaintiff was "a party" who had prepared the materials in anticipation of litigation. Further, the court held that the plaintiff did not waive that protection and her expectation of confidentiality by using ChatGPT. "[T]he work-product waiver has to be a waiver to an adversary or in a way likely to get in an adversary's

© 2026 Thomson Reuters

hand," the court wrote—which was not the case based on use of AI. The court stated that, if the defendants' waiver theory were accepted, it "would nullify work-product protection in nearly every modern drafting environment, a result no court has endorsed."

The court stressed that generative AI programs "are *tools, not persons*." The court criticized the defendant for engaging in a "fishing expedition" that effectively sought to discover the plaintiff's "internal analysis and mental impressions (*i.e.*, her thought process) rather than any existing document or evidence."

***Morgan*—AI interactions were protected, but the tool used had to be disclosed and the protective order could expressly permit only use of AI platforms with confidentiality protections.**

In *Morgan,* a Colorado federal court held that a *pro se* plaintiff was protected under the work product doctrine from producing his interactions with an AI platform—but that he had to disclose *which* AI platform he used to process confidential information, as he had not established that his identifying the tool he used would reveal his mental impressions and case strategy. The court also addressed how a protective order should be amended to address both parties' use of AI tools.

The plaintiff claimed that he was subjected to a hostile work environment and then was terminated based on his race and national origin and in retaliation for whistleblowing activity. The court concluded that the work product rule, as provided in the Federal Rules of Civil Procedure, applies to *pro se* litigants and extends to the use of AI tools. The Federal Rules protect "material prepared in anticipation of litigation or for trial by or for another party," and, under a plain reading, they do not condition protection on the involvement of counsel, the court stressed. The court noted that this conclusion was appropriate especially for a *pro se* litigant: "[*P*]*ro se* litigants are forced to act as both party and advocate, simultaneously. And for the first time in history, widespread access to powerful technology may make that dual role surmountable," the court wrote.

The court concluded that the work product protection is not waived by "sharing" information with an AI platform, just as, for example, emails are not protected although they are shared with an email server. Although AI use technically "discloses" information to a third party, it is not in a manner that makes it likely it will fall into the hands of an adversary, the court stated. "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."

The court distinguished *Heppner* on the basis that case was a criminal matter and thus not governed by the Federal Rules of Civil Procedure; and that it involved a litigant who was represented by counsel but acted entirely apart from his attorney, while this case involved a *pro se* litigant serving as his own legal advocate.

The court also addressed how the existing protective order, which generally prohibited the disclosure of confidential information to third parties without the other litigant's consent, should be amended to address the parties' use of AI tools. The court approved protective order language that prohibited both parties from inputting confidential information into any modern AI platform (including generative, analytical or large language model-based tools) unless the terms of use contractually prohibited the provider from storing or using inputs for training and from disclosing inputs to third parties (except where essential to service delivery); and allows for the ability to delete information on request. The court acknowledged that these requirements would effectively "bar the parties from using most, if not all, mainstream low-to-no-cost AI" for confidential information.

***Tate*—Materials uploaded to AI were protected with respect to their content, but the fact and scope of AI usage had to be disclosed**.

In *Tate*, the Texas Business Court addressed a dispute that arose between litigants during trial when the defendants learned that, throughout the litigation, the plaintiff's principal (a non-lawyer) had been uploading case-related materials to ChatGPT. The defendants moved to compel discovery of the uploaded materials. Citing *Heppner*, the

© 2026 Thomson Reuters

defendants contended that a non-lawyer's exchanges with a generative AI tool carry no expectation of privacy; and that, even if they did, the protection would be waived by sharing the materials with the AI tool. The plaintiff asserted work product protection under Texas Rule of Civil Procedure 192.5, which provides protection to "material prepared or mental impressions developed in anticipation of litigation or for trial by or for a party"—language that is broader than the federal rule applied in *Heppner*.

The court agreed with the plaintiff that the materials were not discoverable based on Texas' version of the work product doctrine as articulated in Texas Rule 192.5. The court, citing *Warner* and *Morgan*, held that the protection was not waived, as uploading the material to ChatGPT did not constitute disclosure of the materials to an adversary or in a manner likely to reach an adversary. The court viewed the sharing of discovery materials with ChatGPT as analogous to using "Westlaw, LexisNexis, e-discovery platforms, or other litigation-support tools."

While the court did not order the plaintiff to produce and disclose the content of the materials that were uploaded to ChatGPT, it ordered that the plaintiff identify, by Bates number, every discovery document that it had shared with ChatGPT, including materials that had been designated as "Confidential" under a protective order. The Court also recommended that the parties amend the protective order to address AI tool usage going forward.

PRACTICE POINTS

Companies should:

- Monitor the evolution of the law in this area as more cases are decided;
- Assume that AI usage may be discoverable, at least to some extent, in litigation;
- Implement policies and protocols with respect to AI usage for legal advice in order to minimize the risk of discovery of AI queries and outputs;
- Inform and train directors, management, and employees with respect to AI use on legal-related matters and issues;
- Consider requiring that legal counsel direct any use of generative AI in connection with seeking legal advice, or at least in connection with pending or anticipated litigation;
- Consider implementing other practices that could mitigate the risk of discovery—such as restricting use of AI for legal advice to specified platforms with favorable terms of use with respect to confidentiality; prohibiting employees other than in-house lawyers from seeking legal advice from AI on behalf of the company, and prohibiting the sharing of the AI interactions with non-lawyers; and, if possible, marking AI interactions as "privileged and confidential" or "prepared by or at the behest of counsel";
- Review existing AI platforms' terms of use to determine whether the confidentiality and related provisions (such as disclaimers with respect to providing legal advice) are likely to be protective or not; and review the terms of use of new AI platforms that the company considers using;
- Consider prohibiting, restricting or requiring disclosure of any personal AI accounts (including on their personal devices) used by directors, officers, or employees to seek legal advice on company-related matters;
- Expressly address generative AI interactions (queries, responses, and AI-generated drafts) in the company's regular document retention policies, discovery requests, litigation holds, preservation obligations, and protective orders; and
- Remain sensitive to the timing of a litigation hold memo (as the obligation to preserve evidence begins when litigation becomes "reasonably foreseeable"), and carefully consider the appropriate instructions to include in the memo (including with respect to preserving chats, prompts, and AI output; disabling auto-delete, temporary chat and incognito

© 2026 Thomson Reuters

features; disclosing shared workspace or other factors that permit others to see prompts or output; and preserving custom or saved settings and instructions to the AI tool).

ENDNOTES:

[1]*United States v. Heppner*, 820 F. Supp. 3d 292 (S.D. N.Y. 2026).

[2]*Fortis Advisors, LLC v. Krafton, Inc.*, 354 A.3d 906 (Del. Ch. 2026).

[3]*Warner v. Gilbarco, Inc.*, 820 F. Supp. 3d 629 (E.D. Mich. 2026).

[4]*Morgan v. V2X, Inc.*, 2026 WL 864223 (D. Colo. 2026).

[5]*Tate Group Automotive v. Legacy Automotive Capital* (Tex. Bus. Ct. June 4, 2026).

BANK-LIKE CUSTOMER IDENTIFICATION PROGRAM RULES PROPOSED FOR STABLECOIN ISSUERS

By Stephen T. Gannon and Michael Treves

Steve Gannon is a partner, and Michael Treves is an associate, in the Washington D.C. office of Davis Wright Tremaine LLP. Contact: stevegannon@dwt.com or michaeltreves@dwt.com.

The build-out of a comprehensive regulatory framework for payment stablecoins is steadily moving toward completion, as the various agencies charged with implementing the GENIUS Act continue to fill in the remaining gaps. The latest piece of that framework is a joint notice of proposed rulemaking[1] by the Financial Crimes Enforcement Network ("FinCEN"), Office of the Comptroller of the Currency ("OCC"), Federal Reserve Board ("Fed"), Federal Deposit Insurance Corporation ("FDIC"), and National Credit Union Administration ("NCUA") governing how permitted payment stablecoin issuers ("PPSIs") must collect and verify identifying information of their customers. The proposal implements the GENIUS Act's directive that PPSIs be treated as financial institutions under the Bank Secrecy Act ("BSA") and maintain an effective customer identification program ("CIP"). Comments on the proposed rulemaking are due August 21, 2026.

The proposal slots into the broader architecture currently advanced through proposals and standards addressing stablecoin reserve, capital, liquidity, and risk management requirements, as well as a separate rulemaking applying anti-money laundering/countering the financing of terrorism ("AML/CFT") and sanctions compliance program obligations to PPSIs. The degree of coordination among the federal agencies in building out the payment stablecoin regulatory framework is remarkable and a notable shift from a previous, often siloed approach. Joint rules applied uniformly across each agency's supervised institutions makes rulemakings more understandable and ultimately compliance less difficult.

KEY TAKEAWAYS

- **PPSIs to be treated as "financial institutions."** The GENIUS Act and the proposed rulemaking treat PPSIs as financial institutions for BSA purposes, subjecting them to the same core AML/CFT program and customer identification mandates that apply to banks.
- **Mandatory written, risk-based CIP.** Each PPSI will have to establish and maintain a written CIP, tailored to its size and business and embedded within its broader AML/CFT program, with risk-based procedures to form a reasonable belief that it knows the true identity of each customer.
- **Core identifying information mirrors bank model.** Before opening an account, a PPSI will have to collect, at a minimum, a customer's name; date of birth (for individuals) or date of formation (for entities); a physical address; and an identification number.
- **Definition of "account" tailored to stablecoin activity.** The proposed rules would expand the concept of an "account" to capture primary market

© 2026 Thomson Reuters