The Chatbot in the Courtroom: What U.S. v. Heppner Means for Health Tec Founders Who Use AI to Think Throug Legal Problems

FEB 15, 2026 • PAID

Abstract

The February 10, 2026 bench ruling in U.S. v. Heppner (No. 25-cr-00503-JSR, S.I is the first major federal decision holding that a client's use of a consumer AI too prepare defense strategy documents waived attorney-client privilege. Judge Jed Rakoff's reasoning rested on three pillars: (1) AI tools are not attorneys, (2) consu ToS explicitly disclaim confidentiality and permit training on inputs and disclos gov't authorities, and (3) pre-existing non-privileged documents don't become privileged just by being sent to a lawyer after the fact. This matters disproportio to health tech founders and executives who routinely use consumer AI to draft, summarize, or analyze content that touches on FDA regulatory submissions, HIP compliance memos, investor deal terms, IP licensing strategy, employment dispu board communications, and reimbursement appeals. The ruling does not automatically doom enterprise AI use – courts are likely to treat tools that contractually prohibit training on inputs and maintain data confidentiality diffe – but the line between "safe" and "not safe" is less obvious than most people thi

Key takeaways:

- Consumer AI (ChatGPT free tier, Claude.ai free/Pro) = no pri protection per Heppner

- Enterprise/API tiers with zero-data-retention = likely safer, not yet definitively tested in court

- AI note-takers on calls with counsel = almost certainly problematic

- Work product doctrine is a separate shield, and it also failed in Heppner

- The three-part test for privilege: attorney, confidentiality, legal-advice purpose consumer AI fails all three

- Heppner had Quinn Emanuel and still lost – this isn't just a DIY legal-research problem

- Health tech-specific exposure: HIPAA enforcement, FDA 483s, SEC/FTC investigations, M&A due diligence, cap table disputes

Table of Contents

What Actually Happened in Heppner

Before getting into implications, it helps to actually understand the facts, becau lot of the commentary floating around gets them a little wrong in ways that matt

Bradley Heppner was the former CEO of Beneficient, an alternative finance com in Dallas. He got indicted in October 2025 on securities fraud, wire fraud, consp and making false statements to auditors in connection with an alleged scheme th prosecutors say cost GWG Holdings investors roughly $1 billion – GWG filed fo bankruptcy, and prosecutors allege Heppner extracted over $150 million for him before that happened, including $40 million to renovate his Dallas mansion. He represented by Quinn Emanuel, which is as serious a white-collar defense firm a can hire.

Here's the key sequence. After Heppner knew he was a law enforcement target a had retained counsel, he used the consumer version of a commercial AI tool to r queries related to the government's investigation. He fed information he'd learne from his Quinn Emanuel attorneys into the AI, generated 31 documents of prom and AI responses, and then transmitted those documents to his lawyers. He appa thought this would help him organize his thinking and prepare for the case. Wh FBI executed a search warrant on his home, they seized devices containing those documents. His legal team asserted both attorney-client privilege and work prod protection. The government moved to compel production. Judge Rakoff ruled fr the bench on Feb. 10, 2026 that neither doctrine protected the documents, sayin not seeing remotely any basis for any claim of attorney-client privilege."

A few things worth flagging that people tend to glide past. First, Heppner wasn' some naive founder who didn't know he was in legal jeopardy – he was under act investigation, knew it, and had already engaged counsel. Second, he had actually received substantive legal strategy input from Quinn Emanuel attorneys and the incorporated that information into his AI queries. Third, his lawyers argued strenuously that the documents should be protected, and they didn't win. If you imagining that "well, I was just using AI to prep some notes before calling my la puts you in a different category than Heppner, think that through more carefully

The Three Ways Privilege Dies

Attorney-client privilege in federal courts has three required elements. You nee communication, made between only privileged parties (meaning attorney and cli made for the purpose of obtaining or providing legal advice. Lose any one of tho three and the communication is not privileged. In Heppner, the government argu and Judge Rakoff agreed – that the AI documents failed on all three.

The most intuitive problem is the "privileged parties" element. A consumer AI t not an attorney. It has no law license, no bar admission, no professional responsi obligations, no duty of loyalty to the client, and no fiduciary relationship with ar It cannot form an attorney-client relationship by definition. The government cit 2025 ruling, In re OpenAI, Inc., Copyright Infringement Litig., for the propositi that "the discussion of legal issues between two non-attorneys is not protected b attorney-client privilege." That's a blunt but accurate statement of existing law. Heppner was typing queries into an AI interface, he was not communicating wit attorney. He was communicating with a piece of software operated by a third-pa company.

The confidentiality problem is the more strategically interesting one, and it has most direct implications for everyday founder behavior. Privilege is lost when information is shared with third parties outside the attorney-client relationship. consumer version of the AI tool Heppner used – and this applies to essentially al consumer-tier AI products – has a privacy policy that explicitly states inputs ma used to train the model and may be disclosed to governmental regulatory author and third parties. Judge Rakoff was direct: the defendant "disclosed it to a third in effect, AI, which had an express provision that what was submitted was not confidential." That's not a technicality. That's the court reading the terms of serv nobody reads and applying them as written.

The third failure point – the legal-advice purpose – is more subtle. The terms of service for consumer AI products typically include explicit disclaimers that the t does not provide legal advice and that users should consult qualified attorneys fo legal matters. Courts have been attentive to those disclaimers. If the platform its telling you it is not providing legal advice, a court is going to be skeptical of you claim that you were using it to obtain legal advice within a privileged relationsh

There's an important practical note here about the third-party disclosure logic t applies to founders beyond formal legal disputes. In health tech, companies regu use AI to summarize regulatory submissions, draft responses to FDA 483 observations, analyze reimbursement policies, or review HIPAA compliance frameworks. If a consumer AI tool is processing that information and retaining r to use it or disclose it to authorities, the privilege analysis is the same even if the no active investigation. The issue isn't just whether you end up in litigation – it's the moment you're in litigation, those documents become discoverable, and the counterparty can use them against you.

Work Product Doctrine: The Second Lir of Defense That Also Collapsed

Some founders reading about Heppner are thinking "okay, but surely work prod doctrine covers this." It's worth spending time on why that reasoning doesn't hol because it's a genuinely common misunderstanding.

The work product doctrine, stemming from Hickman v. Taylor and codified in F Rule of Civil Procedure 26(b)(3), protects materials prepared by or at the directio counsel in anticipation of litigation. It's a different shield than attorney-client privilege – designed to protect an attorney's thought process, strategy, and ment impressions from discovery by opposing counsel. It's also not absolute. Ordinary product can be overcome by a showing of substantial need. Opinion work produ which reflects an attorney's mental impressions and legal theories, gets much str protection.

In Heppner, the work product argument failed for a simple reason that Quinn Emanuel actually conceded: Heppner created the documents "of his own volition and defense counsel "did not direct him" to run the AI queries. That concession dispositive. Work product protection doesn't attach to a client's independent res project, even if the client incorporates information they learned from their lawy attaches to materials prepared by attorneys or at attorneys' explicit direction. A

who decides on their own to run legal questions through a chatbot has not been directed by counsel to do that work. The doctrine simply doesn't reach that far.

The government made a useful analogy the court found persuasive. If Heppner h done Google searches about his legal situation, or gone to a library and checked books about securities fraud law, those search histories and library records woul become privileged just because he later discussed what he learned with his lawy The same logic applies to AI queries. The underlying research doesn't transform protected work product by virtue of being transmitted to an attorney after the fa

Judge Rakoff flagged a wrinkle the government may not have fully anticipated. Because the AI documents incorporated information that Quinn Emanuel had conveyed to Heppner – legal strategy details he'd learned from his lawyers – adm those documents at trial could create a situation where the defense attorneys wo called as witnesses to testify about what they told their client. That witness-advo conflict could cause substantial complications, potentially including a mistrial. I put the government on notice about that dynamic. For founders, the lesson is tha messiness doesn't protect you – it just makes everything more expensive and unpredictable.

Why Health Tech Is More Exposed Thar Most Sectors

Most industries have legal exposure that is episodic. Health tech's legal exposure structural. The regulatory environment is so dense, and the consequences of regulatory missteps so severe, that founders and executives in this space are continuously working through legal questions as part of normal operations. Tha fundamentally different from a SaaS company that deals with legal issues mainly during fundraising, M&A, and employment disputes.

Think about what a health tech executive actually uses AI for in the course of a normal quarter. FDA 510(k) summaries and 513(g) requests involve legal analysis touching on predicate devices, indications for use, and classification questions – with significant legal implications. HIPAA risk assessments and BAA analysis re

parsing regulatory text in ways that are exactly the kind of task AI is good at. Reimbursement and coverage determination appeals involve CMS regulatory frameworks that are genuinely complex and where counsel is often involved. Pay contract analysis – understanding carve-outs, exclusions, termination rights – is something founders routinely offload to AI. Cap table analysis, SAFEs, convertib note terms, MFN provisions – this is due diligence content that regularly involve counsel. FDA inspection responses, particularly to 483 observations, are often dr collaboratively with outside counsel.

In all of these contexts, if a founder or exec is running content through a consun tool – even just to get a plain-English summary of something their lawyer sent th they may be waiving privilege over those communications. And in health tech specifically, the regulatory agencies most likely to be investigating you (FDA, OC HIPAA, CMS, FTC, DOJ, SEC for public or pre-IPO companies) are the same governmental authorities that consumer AI privacy policies explicitly name as pa to whom information may be disclosed. That's not a hypothetical risk. The Hepp privacy policy language that sunk his privilege claim referenced disclosure to "governmental regulatory authorities." The OCR, FDA, and SEC are government regulatory authorities.

The investor dimension matters too. Health tech companies, particularly those in IPO stages, regularly have communications between executives and counsel abou M&A strategy, competitive positioning, cap table mechanics, and investor rights privileged legal strategy is being processed through consumer AI tools – even fo summarization – the Heppner reasoning applies. An acquirer's counsel during d diligence, or a plaintiff's attorney in a shareholder dispute, now has a roadmap f arguing those communications are discoverable. The FTC, which has dramatical ramped up enforcement in digital health post-GoodRx and BetterHelp, makes th even more acute. Companies increasingly likely to face government inquiries at s point in their lifecycle need to treat the Heppner ruling as a near-term operation concern, not a theoretical one.

The Enterprise Carve-Out: Real Safety Just a Better Story to Tell Your Board

The Debevoise analysis of Heppner flagged the most important practical questio does the same analysis apply to enterprise AI tools that contractually prohibit tr on inputs and maintain confidentiality of those inputs? The honest answer is "probably not, but it hasn't been definitively decided."

The logic is straightforward. The court's reasoning rested heavily on the consum tool's explicit terms permitting training on inputs and disclosure to third partie including government authorities. An enterprise agreement that contractually prohibits both of those things removes the factual predicate for the third-party disclosure argument. If the AI provider has committed in writing not to train on data, not to share it with third parties, and to maintain confidentiality of inputs, the "voluntary disclosure to a third party" analysis looks different. The AI tool w function more like a piece of software your law firm uses internally – a means of processing information rather than a disclosure of that information to an indepe party with its own rights to use or share it.

There are caveats founders should not hand-wave away. First, "enterprise" is not magic word. You need to actually read the contract. Some products that market themselves as enterprise tools still reserve significant rights over input data. Zer data-retention policies vary in their scope and their exceptions. Some enterprise exclude certain types of inputs from confidentiality protections, or retain data fo security monitoring purposes with broad enough language to potentially create disclosure risks. The label matters less than the actual contractual terms. Secon even if the enterprise tool maintains input confidentiality, the "attorney-client relationship" problem doesn't fully disappear. A court applying Heppner's reaso could still note that the AI tool is not an attorney, owes no duty of loyalty, and ca form a professional relationship. Third, the "prepared by or at the direction of counsel" requirement for work product protection applies regardless of which A you're using. If you're running your own legal questions through any AI tool wit

your lawyer explicitly directing you to do so, you've not satisfied the work produ
doctrine.

Enterprise AI with contractually sound zero-data-retention terms is likely meaningfully safer than consumer AI for legally sensitive work. But "safer" isn't "safe," and the difference between enterprise and consumer tiers is only one vari in a multi-variable analysis.

The AI Note-Taker Problem Nobody Is Talking About

If consumer AI chatbots for legal research are the visible part of the Heppner ic AI meeting note-takers are the part that's going to hit a lot more companies whe they're not looking. The note-taker risk is, if anything, more severe than the cha risk, because it's more passive and more pervasive.

Tools like Otter.ai, Fireflies, Grain, Fathom, and dozens of othe automatically join calendar-linked calls and produce transcripts and summaries. They've become so frictionless that many founders and executives have them rur by default on essentially every meeting. The problem is that "every meeting" inc calls where attorney-client privilege is being asserted or created. Board calls wit general counsel, strategy discussions with outside counsel about regulatory submissions, M&A calls where lawyers are advising on deal structure – all of the become problematic if an AI note-taker is running on the call.

The Heppner analysis applies with full force here. Consumer-tier note-taking se have their own data retention policies and terms of service. Most retain transcri process them through AI models, and have terms that permit various forms of d use. When a call that would otherwise be privileged is simultaneously being pro by a third-party note-taking service, the voluntary disclosure to a third party pr is immediate and clear. You don't even need to affirmatively send anything to a l afterward – the note-taker is already the third-party disclosure that destroys confidentiality.

There are additional complications specific to healthcare. Many of these calls inv not just attorney-client communications but also HIPAA-protected health information, commercially sensitive data about patients or clinical outcomes, an research data that may be subject to separate confidentiality obligations. The AI taker sitting in on a call between a health tech executive and outside counsel abc say, a potential HIPAA enforcement inquiry is simultaneously destroying privile potentially processing PHI through a third-party system without a BAA, and cre a retention record that could be subpoenaed. That's a lot of legal exposure per meeting transcript. The fix here is operationally simple but requires actually implementing and enforcing a policy. Calls with legal counsel should have a star protocol that no AI note-takers are running. Enterprise note-taking solutions wi robust BAAs and contractual data-use restrictions can potentially be used in sor contexts, but the same caveats that apply to enterprise AI chatbots apply here.

Practical Protocol for Founders and Executives Right Now

The Heppner ruling doesn't mean founders have to stop using AI for legal-adjac work entirely. It means being intentional about where the line is, and actually building operational hygiene around that line rather than assuming it's fine beca nothing bad has happened yet.

The highest-risk behavior right now is using any consumer-tier AI tool – free or personal tiers of any major AI product – to process content that has already beer communicated to you by counsel or that you intend to share with counsel as part active legal matter. That's the Heppner fact pattern almost exactly. If you're und regulatory scrutiny, in a funding dispute, navigating an employment claim, or dc anything involving an active attorney-client relationship and privileged legal str do not run any of that through consumer AI. The privilege waiver is retroactive a extends to the underlying communications that the AI processed, not just the Al output itself.

The next tier of risk is using consumer AI for legal-adjacent research or analysis areas where you reasonably expect future legal exposure. In health tech, that's al all of the terrain described above. FDA strategy, HIPAA compliance, reimbursen policy, IP licensing, investor rights – all areas where consumer AI processing cre waiver risk that may only manifest later. The better operational approach is to u enterprise AI tools with proper data governance, or to route those questions dire through counsel without AI intermediation.

For investor-grade diligence, the implication is that health tech companies shou AI governance to the legal and compliance sections of their data rooms. Sophisti acquirers and late-stage investors are going to start asking whether executives h used consumer AI tools to process privileged communications, particularly in regulated companies where the regulatory history is a material part of the transa value. Having a policy in place – and a documented practice of using compliant t is basic risk hygiene at this point.

The attorney-direction point from the work product analysis also has a useful pc implication. If counsel explicitly directs a client to use AI tools for specific resea analysis tasks, that direction creates a stronger argument for work product prote Founders who want to use AI as part of their legal workflow should discuss that practice explicitly with their lawyers, document the direction, and ensure the to being used are on an approved list that counsel has blessed. That's not a foolpro shield, but it's meaningfully different from the Heppner situation where the defendant acted entirely on his own initiative.

Enterprise AI adoption in the legal workflow is not inherently problematic, but requires actual procurement discipline. Zero-data-retention isn't a checkbox – it contractual term that needs to be negotiated, documented, and periodically verif The same way a health tech company would diligence a cloud vendor's HIPAA compliance before routing PHI through them, they need to diligence an AI vend data governance terms before routing privileged communications through them. current practice in most companies is essentially ad hoc. Post-Heppner, that's a liability. The founders who build thoughtful AI governance now – for their prod for their internal operations, and for how they interact with their legal teams – a

going to be in materially better shape than those who don’t when the next ruling arrives. And there will be a next ruling.

2 Likes • 1 Restack

← Previous

Nex

Discussion about this post

Comments Restacks

Write a comment...