Satya's practice focuses on technology transactions, intellectual property, and commercial agreements for startup and growth-stage technology companies across AI, robotics, cloud computing, and emerging tech sectors.

Email: snarayan@gcalaw.com
LinkedIn Newsletter: Techline Brief
Connect on LinkedIn

Satya S. Narayan
June 2, 2026
Subscribe to my newsletter: Techline Brief

When You Suspect a Legal Problem, Call Your Lawyer - Not Claude or ChatGPT

Over the next several weeks, I am focusing on one of the most underappreciated legal risks of the AI era: **discoverability**. Each post will examine a different AI tool and how the records it generates can become a serious liability in litigation.

I will start with a case that should be required reading for anyone who has ever typed a sensitive question that might have legal implications into a generative AI system.

The Case: United States v. Heppner

In February 2026, the Southern District of New York decided *United States v. Heppner*, an early and important ruling addressing AI use and attorney-client privilege. Bradley Heppner, a former executive facing securities and wire fraud charges, had already received a grand jury subpoena and retained counsel. Without any suggestion from his lawyers that he do so, he used Anthropic's Claude to draft defense strategy documents and outline responses to potential charges.

The FBI later seized approximately 31 of those documents during a search of his home. His attorneys asserted privilege, specifically the attorney-client privilege and the work product doctrine. The court rejected the claim.

Judge Rakoff ruled on several grounds. Two are particularly relevant here. First, Claude is not a lawyer; it is not licensed, owes no fiduciary duty, and is not bound by attorney-client confidentiality. Second, and most importantly for the broader lesson, Heppner acted entirely on his own initiative, without any direction from counsel. The court's other grounds, including its analysis of Anthropic's privacy policy and the role of third-party platforms, have attracted commentary and debate in the legal community, but those questions do not change the practical lesson this case highlights.

Heppner also made matters worse by inputting information his own counsel had shared with him, effectively waiving whatever privilege may have attached to that information. As Judge Rakoff put it, non-privileged communications are not “alchemically changed into privileged ones” simply because they are later shared with counsel. The damage was done at the moment of creation.

Why the “without direction of counsel” point matters most

The attorney-client privilege and work product doctrine are not general shields against disclosure of anything a person writes while thinking about a legal problem. They protect specific communications between a client and attorney, made in confidence, for the purpose of obtaining legal advice, and work product prepared by or at the direction of counsel in anticipation of litigation.

Heppner's counsel conceded at oral argument that they had not directed him to use Claude. That concession was fatal. Had counsel directed Heppner to use the tool, Judge Rakoff acknowledged that 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." The court referenced in passing two cases: *United States v. Adlman*, 68 F.3d 1495 (2d Cir. 1995), which involved an attorney acting in a dual legal and business advisory capacity, and *United States v. Kovel*, 296 F.2d 918 (2d Cir. 1961), which extended privilege to a tax accountant acting as counsel's agent under appropriate supervision. Both cases were decided in the context of human professionals. Whether courts will extend that reasoning to a generative AI tool, affording it the same treatment as a human professional agent of counsel, remains an open question. But because Heppner acted on his own, the documents were simply not privileged to begin with, and no subsequent sharing with lawyers could fix that.

This is the point the case most clearly establishes, and it applies far beyond criminal defendants.

Any time a person independently uses a generative AI tool to work through a problem that could have legal implications, without being directed by counsel to do so, the entire exchange, the input and output, starts life as an unprotected document.

To see why this matters in practice, consider a more common scenario:

An executive senses a potential problem: a contract dispute, a regulatory issue, a personnel matter with legal risk. Before calling counsel, they open ChatGPT or Claude to "think it through."

That instinct is understandable. It is also risky.

The real risk: creating discoverable records you would never otherwise write

The risk is not limited to situations where a legal problem is already visible. Ordinary business decisions, pricing strategies, competitive responses, acquisition rationales, and personnel matters

can become legally significant long after the fact, as decades of antitrust and products liability litigation have shown.

In either case, every prompt entered into the generative AI system is a record. Every response generated is a record. Under the Federal Rules of Civil Procedure, that entire input/output exchange qualifies as electronically stored information (ESI), discoverable in litigation just like email, Slack messages, or text messages. And unlike those formats, AI chat records may also exist on the platform's servers, reachable by subpoena to the company, not just by searching the user's own devices. Deleting your local copy does not necessarily solve the problem.

But generative AI records are a different kind of ESI altogether.

People tend to be far more candid with AI tools than they are in communications with other humans. They lay out facts more completely, test worst-case theories, ask questions they would hesitate to put in an email, and iterate and expand their thinking in writing.

The result is not just a record, but what can amount to a highly structured, unusually candid narrative with potential for legal exposure, often created before any lawyer is involved and outside any protection that lawyer involvement might have provided.

That is what makes these exchanges so valuable in discovery.

This is not about banning AI. It is about recognizing when and how to use it.

None of this means generative AI tools are inherently inappropriate. Using them for general education, understanding general concepts, or non-sensitive matters is a different question.

The risk emerges when the issue becomes fact-specific and potentially consequential, especially in a legal sense, whether or not that consequence is yet visible. That includes situations where litigation or regulatory scrutiny is already suspected, but also ordinary business decisions that could become legally significant in hindsight.

At that point, the question is no longer "What is the law?" It is "How does the law apply to these specific facts?"

That is the moment to call counsel first.

The practical takeaway

This is not complicated.

If you think you may have a legal issue or you are working through a business decision that could become legally significant, resist the instinct to quietly work through it in a generative AI tool first. That initial thinking-out-loud may become a discoverable record later, one that is more detailed and more candid than anything you would have written to another person. It may have

been created before any legal privilege attached, and it may be stored in places beyond your control.

The same principle applies to employees at every level. If a situation looks like it might have legal implications for the company, the first call should be to counsel, not to Claude, ChatGPT, or any other generative AI tool.

The conversation with legal should include whether to use a generative AI tool in connection with the matter at all. If counsel determines that AI tools are appropriate, get clear legal guidance on how and for what purposes the tool may be used in connection with the matter. Any tool used should be enterprise-grade, meaning the company has worked with the AI vendor to ensure confidential exchanges, no training of the model on company inputs or outputs, no unauthorized disclosure, and clear data retention limits.

The threshold requirement in all of this is the same. Whether you suspect a legal problem or are working through a business decision that could become one, call your lawyer before you call on AI.

More on discoverability in the next few weeks.

Sources

- *United States v. Heppner, No. 25 Cr. 503 (JSR), 2026 BL 52143 (S.D.N.Y. Feb. 17, 2026).*
- *United States v. Adlman, 68 F.3d 1495 (2d Cir. 1995).*
- *United States v. Kovel, 296 F.2d 918 (2d Cir. 1961).*
