

Me: I'm concerned about attorney-client privilege regarding giving you client-identifying details. Is my paid individual Claude.ai plan sufficient to keep that privilege if I disclose those details to you?

Claude.ai:

This is exactly the right question to ask, and I want to give you a completely straight answer: **I would not rely on your paid Claude.ai subscription as sufficient protection for attorney-client privilege, and I'd recommend against sharing client-identifying details with me.** Here's why:

The Core Problem: Third-Party Disclosure Risk

Attorney-client privilege can be waived when confidential communications are disclosed to third parties outside the privilege. Whether an AI tool like Claude qualifies as a protected "agent" of the attorney – analogous to a paralegal or associate – is **genuinely unsettled law**. No Illinois court has ruled on it, and the ISBA has not issued a definitive ethics opinion specifically on this point that I can cite with confidence.

The Confidentiality Obligation Is Separate and Also at Risk

Even apart from privilege, **Illinois Rule of Professional Conduct 1.6** requires reasonable measures to prevent unauthorized disclosure of client information. The question of whether using a commercial AI service satisfies "reasonable measures" is unresolved. Comment [18] to Rule 1.6 addresses technology generally, but predates generative AI tools.

What the ABA Has Said

ABA Formal Opinion 512 (2024) – which is real and I'd encourage you to read it – directly addresses generative AI use by lawyers. It concludes that before inputting client confidential information into an AI tool, a lawyer must:

- Understand how the AI provider uses, stores, and potentially trains on submitted data
- Determine whether the provider's terms adequately protect confidentiality
- Take steps to avoid inadvertent disclosure

What Anthropic's Paid Plan Actually Provides

I'd encourage you to read Anthropic's current privacy policy and terms of service directly rather than rely on my characterization – those terms can change, and **you** need to make that assessment, not me. The key questions to answer from those documents are:

1. Does Anthropic use conversation content to train future models?
2. Who can access conversation data, and under what circumstances?
3. Is there a data processing agreement (DPA) available for professional/legal use?
4. What happens if Anthropic receives a subpoena or legal process for your data?

Anthropic does offer enterprise agreements with stronger data protections – that tier may be more appropriate for legal work involving client information.

My Practical Recommendation

****For our work together, the safest approach is what we've already been doing – use placeholders.**** A letter I draft is just as effective with [CLIENT FULL NAME] and [GAME TITLE] as it would be with real names. You fill those in on your end, in your protected environment.

For the specific facts I actually **need** to do good legal work – the nature of the IP, the procedural history, the legal issues – those can almost always be described in general terms without identifying the client.

Bottom Line

Don't share client-identifying details with me until you've:

1. Read ABA Formal Opinion 512
2. Reviewed Anthropic's current terms and privacy policy yourself
3. Potentially consulted the ISBA Ethics Hotline (they're responsive and this is exactly what it's for)
4. Decided whether an enterprise-tier agreement with a data processing addendum is warranted for your practice

I'd rather slow you down slightly than have you inadvertently create an ethics exposure. The placeholder approach costs us nothing and protects everything.

Me: Placeholder is good for now. I do have training turned off, but I'm still wary, unless that changes things dramatically.

Claude.ai:

On Training Being Turned Off

It helps, but it doesn't change my recommendation dramatically. Here's why: Having training turned off addresses one of the concerns from ABA Opinion 512 – that your data could be used to train the model and thus become more broadly accessible. That's meaningful. However, it doesn't resolve:

Whether Anthropic employees or systems can access conversations

What happens in response to legal process (subpoenas, court orders)

Whether the "third party disclosure" problem for privilege purposes is solved – a court could still find that transmitting data to a commercial server outside your firm constitutes a disclosure regardless of training settings

So – better, but not a green light for client-identifying details. The placeholder approach remains the cleanest solution at no practical cost to us.