STARTUP LAW

The Hidden Legal Risks of Using AI to Build Your Startup

By Ari Tuchman | Tuchman Law, APC | Los Angeles, CA

Startups are exactly where the temptation to treat AI as legal counsel is strongest — technical founders who are comfortable with AI tools, tight budgets, and a constant stream of documents that need to get done fast: NDAs, advisor agreements, employment offers, SAFEs. It's also exactly where the stakes compound the most. A generic contract mistake at an established company is usually fixable. An early-stage documentation mistake gets baked into your cap table, your IP chain, and your governance structure — and it's a future investor's lawyer, not you, who eventually finds it.

Equity Documents Don't Get a Second Chance

A founder or advisor equity agreement that looks mechanically correct — proper defined terms, a vesting schedule formatted the way vesting schedules are supposed to look — can still miss the specific term your situation actually needed: reverse vesting acceleration triggers, IP assignment tied directly to the grant, or provisions that match your state's specific default rules rather than a generic template. The document doesn't fail to compile the way bad code does; it just quietly doesn't do what you assumed it did. And unlike a software bug, you usually don't find out until a co-founder leaves early, or until a Series A term sheet is on the table and opposing counsel starts asking questions about your cap table.

Missing IP Assignment Is Invisible in the Document Itself

Unassigned intellectual property from a contractor, part-time advisor, or early engineer who was a friend-of-a-founder is one of the most common issues that surfaces in venture diligence — and it's a particularly bad fit for AI drafting tools, because the problem isn't a flaw in any single document. It's an absence: a person who should have been in the assignment chain and wasn't. A fluent, well-formatted IP assignment agreement gives you no signal that someone was left off the list entirely. That gap sits invisible until a diligence team goes through your full contributor history and finds the name that isn't anywhere in your paperwork.

Founder Disputes Can Make Your Chat History Discoverable

This is the danger founders think about least, and it should worry them more. Earlier this year, a federal court ruled that materials a defendant generated by discussing case strategy with a consumer AI chatbot weren't protected by attorney-client privilege — because no attorney-client relationship exists with an AI platform, and standard consumer AI terms of service typically allow the company to review submitted data. The same logic applies directly to a founder working through a brewing co-founder dispute, an investor negotiation, or a potential claim by talking it through with a general-purpose AI tool instead of counsel. If that dispute ever

escalates — a business divorce, an investor lawsuit, a departing co-founder's claim — those conversations may not be protected, and they may become exactly the kind of record the other side gets to read.

Fundraising Diligence Is Where Fabrication-Adjacent Errors Get Expensive

The broader pattern showing up in courts — AI producing citations and authority that look completely legitimate and aren't — has a direct startup analog. If a SAFE, convertible note, or governance document was drafted or "verified" by AI as standard when it actually deviated from market terms or your state's requirements, that doesn't surface until a Series A diligence team reads it closely, the same way a fabricated case citation doesn't surface until a judge or opposing counsel checks it. A federal appeals court sanctioned two attorneys this year for exactly this pattern — fabricated authority that looked real until someone was required to verify it — with penalties reaching $15,000 per attorney. Your cap table has to hold up under that same kind of scrutiny before a round closes, and finding the problem during diligence rather than before it means it shows up as a delay, a re-cut round, or a price adjustment.

Regulators Are Watching This From Both Directions

A bill pending in New York would impose liability on companies whose AI chatbots deliver outputs that amount to unauthorized practice of law. That matters if you're using AI as a stand-in for counsel — and it matters again, separately, if your startup is building any product that offers legal-adjacent guidance, since that's precisely the category regulators are starting to target.

What This Means in Practice

AI tools can still speed up legitimate early-stage work — organizing your thinking, producing a first draft, researching general concepts before a real conversation with counsel. The distinction worth holding onto is which documents become permanent. Anything that touches equity, IP assignment, or governance should be routed through actual counsel before signature, not after a problem is found. And anything resembling company strategy or a brewing dispute should be treated as something you'd only otherwise say under privilege — which means not saying it to a consumer AI chatbot at all.

Final Thoughts

The pattern across all of this is the same: AI is good at producing something that looks finished, and a startup's earliest legal documents don't get judged on whether they look finished — they get judged later, by an investor's lawyer or a departing co-founder's lawyer, under conditions you don't control. Routing these decisions through the same counsel consistently from the start, which is exactly what an early-stage general counsel relationship is built for, is what keeps that later scrutiny from finding something you didn't know was there.

About the Author

Ari Tuchman is a Los Angeles transactional attorney and founder of Tuchman Law, APC. He focuses on real estate transactions, business acquisitions, and general counsel services for companies and investors

throughout California.

This post is for general informational purposes only and does not constitute legal advice or establish an attorney-client relationship.