The Bot Said It Was Fine: The Perils of AI-Generated Legal Advice for Dealers and Finance Companies

March 31st, 2026 | Megan Nicholls and Mark D. Metrey

It's a busy day for operations. Customers are everywhere, and business is good. You need to respond to a question about a payment convenience fee. To save time, you quickly enter: "Can I charge an ACH convenience fee in [state]" into your favorite generative artificial intelligence tool. The tool provides you with an instant answer, written confidently and in plain language. The response that takes seconds to create sounds like it was written by a seasoned (and free) compliance attorney. Seems like a perfect tool, right? Well, maybe not.

The problem is that an AI tool is not a lawyer; rather, it is a text prediction system that can be spectacularly helpful for generating ideas and summarizing texts but dangerously unreliable when its output is treated as legal advice. In fact, generative AI tools often prohibit use of their services for legal advice in their terms of use. Responses may sound right, but that does not mean that the responses are right. This is where dealers and finance companies need to beware.

This article highlights the risks of using generative AI and offers practical guardrails for using a generative AI tool to create efficiency without letting it crash you into a legal roadblock.

Using generative AI can be risky

First, generative AI tools, while good at confidently responding to a prompt, are bad at the multi-step processing that is often required prior to answering a compliance question. Answers to compliance questions are usually based on the law, regulations that have implemented the law, case law, and industry practice and often include a contingency based on unknown facts. Generative AI tools have not mastered this process that an attorney might use. For example, in order to answer a simple compliance question, an attorney would typically: (1) conduct research into the law, regulations, cases, etc.; (2) analyze the research; (3) understand the facts of the client's business that are pertinent to answering the question; (4) take into account industry experience; (5) craft a position based on this information; and (6) when applicable, determine and describe the unknown. Instead of one generative AI tool, this process would take several AI tools working together, with each tool performing a singular task. The technology is just not there yet.

Second, the laws, particularly those governing consumer financial services, vary by jurisdiction. For example, state fee caps, disclosures and form requirements, refund timing rules, notice content rules, licensing triggers, and advertising requirements all vary by state. Generative AI is designed to create a summarized response based on the data on which the model is trained. A general response to a compliance question based on a "simple" fact pattern will not capture these jurisdictional variances properly.

Third, ethics authorities have repeatedly warned that generative AI can produce inaccurate output, including hallucinations that look plausible but have no basis in fact. The consistent theme is that users cannot abdicate judgment to the tool and that outputs require independent verification appropriate to the task. Using these responses as authority for a legal or compliance decision creates significant risk, with that risk increasing as the severity of the error(s) and significance of the decision(s) increase.

Fourth, a publicly available generative AI tool may not keep your prompt (or your conversation) confidential. Financial companies hold nonpublic personal information, whether in a system of record for accounts, deal structure, a consumer complaint, a draft demand letter, or an exception request. Blindly pasting this type of information into a generative AI tool risks the security of this information. Even sharing internal pricing, reserve methodology, or litigation strategy with a third party could risk your business's confidential information. For instance, did you know that if you share a link of your digital conversation, the link could now become searchable on the web, potentially exposing what would be privileged or confidential information to countless people?

Practical guardrails to use generative AI without creating additional legal risk

You do not need to ban generative AI to manage the risk. You need guardrails that manage the risk. A common way to document these guardrails is through an AI governance program. Here are some examples of guardrails that you might include in your program:

- Prohibit the use of an AI tool unless the business has evaluated and approved the tool.
- Prohibit the use of an AI tool in high-risk situations, such as for analyzing compliance questions or providing legal advice.
- Prohibit the inclusion of any nonpublic personal information, account details, deal jackets, consumer complaint specifics, or drafts of legal correspondence into non-approved tools.
- Use enterprise (internal) configurations with appropriate security, retention, and contractual terms.
- Train staff. Training should include descriptions of what "confidential" and "privileged" mean and how an AI tool can jeopardize these designations when nonpublic personal information or other confidential information is included in a prompt. It is also a good idea to train staff on prompt drafting generally, as a good prompt will produce a better result.
- Verify all outputs prior to using them for any substantive purpose. Remember, a generative AI tool is really good at sounding confident and conclusive.
- Think of generative AI as a tool to help facilitate creativity or to increase process efficiency by automating a single step in a multi-step process.

Final word

We have yet to see a case where liability was avoided with an excuse of "the bot said it was fine."