

Artificial Intelligence (AI) tools such as **ChatGPT** and **Google Gemini** have transformed how information is accessed, synthesized, and summarized. For family offices and ultra-high-net-worth (UHNW) advisors, these generative models offer new capabilities in preliminary research, parsing lengthy reports, and generating concise summaries. However, these benefits come with a serious caveat: the risk of AI hallucinations.

Understanding AI Hallucinations and Why They Are Problematic

AI hallucinations occur when a language model confidently generates text that is factually inaccurate, fabricated, or unverifiable. Unlike traditional software errors, hallucinations are often *persuasive* and linguistically fluent, making them deceptively credible. For busy professionals who rely on AI tools to expedite complex tasks, this can lead to unwitting dissemination of false information.

In legal contexts, and particularly in scenarios involving legal briefs, court documents, or canadianfamilyoffices.com regulatory filings, the stakes are extremely high. Imagine an AI tool generating a court filing or referencing a legal precedent that doesn't exist, or fabricating citations — what we commonly call *fake citations*. These hallucinations can lead to reputational damage, sanctions, or even legal liabilities.



Why Family Offices Use AI and Face Elevated Risk

Family offices often leverage AI tools like ChatGPT and Google Gemini primarily for:

- Preliminary research on investment opportunities and regulatory environments
- Parsing through dense legal, financial, or market reports
- Generating summaries to expedite decision-making

These use cases, while valuable for efficiency, expose family offices to unique risks because:

- **Private Markets Data Scarcity:** Family offices deal extensively in private equity and alternative asset classes, often relying on scarce, incomplete, or non-public data sources. AI trained mainly on public data may hallucinate details when synthesizing such information.

- **Preliminary vs Final Use:** AI output is often treated as fact rather than a draft or pointer, increasing the risk of relying on inaccurate text without proper verification.
- **Multi-disciplinary Complexity:** Insights often cut across legal, financial, and regulatory domains, increasing the chance that hallucinated references or legal citations slip through.

The Specific Legal Risks of AI Hallucinations

1. Fake Citations and Unsupported Legal References

One of the most dangerous AI hallucinations involves generating legal citations or case law that do not exist. Models like ChatGPT and Google Gemini try to mimic the patterns of legal writing and referencing, but can fabricate docket numbers, case names, or statutes.

For family office teams preparing legal briefs or court filings, even a single fabricated precedent can:

- Undermine the credibility of the document
- Lead to rejection, sanctions, or damage claims if submitted in a legal process
- Expose the office to malpractice or compliance risks

2. Court Filing Risk from Unverified AI Output

When AI-generated drafts are treated as final rather than preliminary, the risk of erroneous filings rises. Courts and regulators expect factual and verifiable documentation. An AI hallucination in a filing could cause procedural setbacks, damage client interests, or even invite legal penalties.

3. Misleading Interpretations in Due Diligence and Regulatory Contexts

Family office teams conducting *manager due diligence* or internal compliance may use AI to summarize regulatory guidance or interpret complex financial reports. AI hallucinations here can lead to:

- Misstating regulatory requirements or exceptions
- Underestimating risks in private market investments
- Suboptimal or erroneous decision-making due to incorrect data extracted or inferred

4. Intellectual Property and Confidentiality Concerns

Hallucinations can also introduce unintentional breaches if AI cites or mimics protected proprietary information or generates text that could be traced to confidential sources. Family offices must be mindful of information governance when using AI tools.

Faithfulness vs Factual Hallucinations: A Critical Distinction

In AI risk management, it's important to distinguish between:

- **Faithful summaries:** AI outputs that closely represent the original data or source material without distortion.
- **Factual hallucinations:** AI-generated inaccuracies that invent or distort facts beyond acceptable paraphrasing.

Many family office teams overlook this nuance, treating AI-generated summaries as fully faithful rather than approximate. Given that private markets data is often fragmented or non-standardized, the risk of hallucinated

facts increases, demanding heightened skepticism and robust validation.



How ChatGPT and Google Gemini Illustrate These Risks

Both ChatGPT and Google Gemini represent advanced Large Language Models (LLMs) designed to interpret, generate, and summarize text using vast language patterns. However, neither is immune from hallucinations:

- **ChatGPT:** Known for fluency and versatility, ChatGPT sometimes hallucinates plausible but false facts or citations, especially in specialized domains like law.
- **Google Gemini:** Touted for deep integration with Google's datasets and real-time information, yet its hallucination risk remains, particularly when dealing with private or scarce data outside public knowledge graphs.

Both tools' default outputs should be treated as *preliminary drafts* requiring human expert verification, particularly for submissions involving legal briefs or official court filings.

Practical Mitigations for Family Offices

To manage legal risks from AI hallucinations, family offices should adopt the following best practices:

1. **Vet AI Outputs Against Verified Sources:** Always cross-check AI-generated citations and claims with authoritative databases and official legal repositories.
2. **Mark AI Outputs as Drafts:** Never treat preliminary AI results as final without rigorous human review, especially in compliance or legal documents.
3. **Invest in AI Literacy and Governance:** Train staff on AI limitations, engage compliance teams early, and incorporate AI governance frameworks that map potential hallucination risks.
4. **Use Specialized Legal AI Tools:** Complement general LLMs like ChatGPT with legal-specific AI platforms equipped with better validation and source traceability.
5. **Document Review Trails:** Maintain audit-ready documentation showing how AI outputs were verified and what human edits were made prior to filing or investment decisions.

6. **Limit AI Use to Appropriate Tasks:** Use generative AI for brainstorming or summarization rather than direct legal filings or final court submissions.

Conclusion

AI hallucinations present substantial legal and reputational risks, especially for family offices operating in the complex, data-sparse environment of private markets. While tools like ChatGPT and Google Gemini unlock tremendous productivity gains in preliminary research and report parsing, they cannot substitute for rigorous human validation.

Understanding the nuances of faithfulness versus hallucination, and incorporating robust governance around AI use in legal briefs and court filings, is essential. Adopting a cautious, auditor-ready stance — always asking “what would we show an auditor?” before finalizing AI-generated documents — can mitigate the risk of costly mistakes and protect client interests over the long term.

References and Further Reading

- [OpenAI ChatGPT Overview](#)
- [Google Gemini Launch](#)
- [Harvard Business Review on AI Risks](#)
- [AI and Legal Ethics – Stanford Law](#)