


Artificial intelligence (AI) tools like **ChatGPT** and **Google Gemini** have rapidly become valuable assets in the family office setting. They assist with preliminary research, parsing complex reports, and generating concise summaries for busy decision-makers. Yet, these AI models are not infallible. One of the most critical risks to understand and communicate effectively is the phenomenon of *AI hallucinations*.

This blog post aims to provide a **plain language briefing** and a **risk framing** approach tailored for boards and family councils. We will cover what AI hallucinations are, why they occur, real examples relevant to family office contexts—especially in private markets—and practical guidance on mitigating this risk.



What Are AI Hallucinations?

AI hallucinations occur when AI language models generate information that is fabricated, inaccurate, or unverifiable but presented confidently and persuasively. Unlike a human making an honest mistake, AI outputs can look highly credible, even to experts.

- **Why the term 'hallucination'?** It's a metaphor borrowed from human neurology, describing when a model "sees" things that aren't there in its training data or knowledge base.
- **How do ChatGPT and Google Gemini factor in?** Both are advanced large language models (LLMs) that generate text responses by predicting likely word sequences based on vast datasets—but they don't "know" facts in a human sense.

Put simply: AI hallucinations = *plausible but false* answers.

Why Is This Risky for Family Offices?

Family offices frequently rely on AI tools for speed and efficiency:

- **Preliminary research:** Quickly gathering background on potential investments, sectors, or geographic areas.
- **Parsing reports:** Extracting key points from lengthy fund documents, auditor reports, or market analyses.
- **Summaries:** Creating executive-level briefs for board materials or family council updates.

However, the very persuasiveness of AI output can lull users into unwarranted confidence. For ultra-high net worth families, where investment decisions involve large capital commitments and regulatory compliance, inaccuracies can cause real financial and reputational harm.

Private Markets Data Scarcity Increases Error Risk

Unlike public markets with abundant, standardised data, private markets suffer from:

- **Limited transparent data:** Fund terms, performance, and valuations are often confidential and less standardized.
- **Small datasets:** AI models have less training data to "learn" private markets context, increasing error likelihood.

This scarcity elevates the chance that AI will "fill in the gaps" with fabrications—classic hallucinations.

Types of AI Hallucinations: Faithfulness vs. Factuality

Type	Description	Example	Implication
Faithfulness Hallucinations	The AI misrepresents or inaccurately summarizes information that actually exists in the source material.	AI summarizes a fund document but incorrectly states the management fee as 1% instead of 1.5%.	Leads to misunderstandings about terms; lowers trust in reporting.
Factual Hallucinations	The AI invents information that does not exist anywhere—pure fabrication.	AI claims a fund has investments in a specific startup that the startup denies and no public record confirms.	False data leads to flawed due diligence or poor investment decisions.

Both types are risky, but factual hallucinations can be more damaging as they may remain undetected until after decisions are made.

Real Examples Relevant to Family Offices

Example 1: Preliminary Research on a Private Equity Fund

Suppose a family office analyst asks ChatGPT:

"Provide key facts about XYZ Private Equity Fund's historical returns and fund size."

The AI returns a confident summary citing 15% net IRR and a fund size of \$2 billion USD.

But in reality, the fund's latest vintage just closed and has no historical returns; the actual fund size is \$1.3 billion CAD. The AI "hallucinates" data from similar funds or outdated sources.



Example 2: Summarizing a Manager Due Diligence Report

Using Google Gemini, a family office team asks for highlights of a report on the fund manager:

"Summarize key investment risks mentioned."

The AI lists regulatory risks and market concentration risk but completely omits the critical risk related to key employee turnover documented in the report.

Because it's a faithfulness hallucination, the AI inaccurately reflects the source, potentially giving the board a false sense of security.

How to Explain This to a Board or Family Council

Boards and family councils are not expected to be tech experts, so the goal is clarity and framing the risks in business terms:

1. **Use Plain Language:** Describe AI hallucinations as "confident but incorrect or made-up answers that machines sometimes give."
2. **Highlight Persuasiveness:** Explain that the AI doesn't "check" facts and can sound very sure by design, increasing risk.
3. **Connect to Business Risks:** Emphasize that errors can lead to poor investment choices or compliance glitches.
4. **Reinforce Human Oversight:** Stress that AI outputs are starting points, not final answers.
5. **Describe Controls:** Illustrate your team's use of data verification, multiple sources, and independent due diligence.

A simple analogy might help:

"Think of AI like a junior analyst who's read many reports but sometimes guesses to fill in blanks—without double-checking. We always verify before we act."

Mitigating AI Hallucination Risks: Practical Tips

Families and their advisors can adopt the following best practices:

- **Ask for Sources:** When using AI like ChatGPT or Google Gemini, explicitly request citations or original sources.
- **Use Multiple Tools:** Cross-check AI outputs across different models and human experts.
- **Verify High-Stakes Data:** Confirm critical financial figures or legal terms against official documents.
- **Train Teams:** Educate analysts and board members about AI's strengths and limitations.
- **Implement Governance:** Establish questions like "What would we show an auditor?" before adopting AI-generated data.
- **Keep Logs:** Document AI queries and results for audit trails and to spot recurring hallucination patterns.

Conclusion

AI tools such as ChatGPT and Google Gemini offer tremendous opportunities to accelerate family office workflows—especially for preliminary research, report parsing, and summarization. Yet, their susceptibility to *hallucinations* demands transparent communication and structured controls when presenting insights to boards or family councils.

By providing a **plain language briefing** that stresses the **risks of persuasive but false output**, real examples from private markets contexts, and a [factual hallucination](#) clear **risk management approach**, family offices can adopt AI confidently without jeopardizing the integrity of critical decisions.

Always remember: AI is a powerful assistant—not a substitute for rigorous human judgment and due diligence.