

In today's high-stakes family office environment, where ultra-high-net-worth (UHNW) capital depends on razor-sharp diligence and informed committee decisions, integrating AI tools like ChatGPT and Google Gemini offers tremendous efficiency gains. Yet, alongside these benefits lurks a critical risk: AI hallucinations. These are subtly fabricated or inaccurate statements presented confidently by AI models—a perilous pitfall for investment committee (IC) decks. This post offers a deep dive into the nature of AI hallucinations, why they pose unique dangers in private markets, and how robust **IC deck controls**, a disciplined **fact checking process**, and the role of a final reviewer are essential to keep decks accurate, trustworthy, and audit-ready.

Understanding AI Hallucinations: The Basics

AI hallucinations occur when language models generate outputs that sound plausible but are factually incorrect or completely fabricated. These are not mere typographical errors, but confident-sounding statements with little to no basis in reality. With tools like ChatGPT and Google Gemini, this risk is magnified given the complexity of investment topics and the high stakes of decision-making.. Pretty simple.

Types of Hallucinations Relevant to Investment Research

- **Factual Hallucinations:** Completely made-up statistics, company facts, or deal terms.
- **Faithfulness Errors:** Misinterpretations or oversimplifications of real data or reports that distort meaning.

Both types can mislead an investment committee, eroding trust and putting client capital at risk.

Why Family Offices Need Extra Vigilance

Family offices often leverage AI primarily to:

- Perform preliminary research on sectors and companies
- Parse and summarize dense industry reports and due diligence notes
- Generate succinct executive summaries for IC deck prep

While this can speed up workflows, it also introduces unique risks:



Private Markets Data Scarcity Amplifies Error Risks

Public market data is relatively abundant and verifiable, but private markets data is often fragmented, sparse, or anecdotal. AI models trained predominantly on public disclosures are more prone to hallucinate details when extrapolating or filling gaps in private equity, venture capital, or direct investments.

- AI may invent financial metrics for a private company that does not disclose them.
- Misstate deal terms or ownership structures with no public record to cross-check.
- Overgeneralize from sector data inappropriately applied to niche family office investments.

Given the paucity of hard data, family offices must adopt an especially stringent *fact checking process* and integrate human expertise to gauge AI outputs' reliability.

Implementing IC Deck Controls to Mitigate Hallucinations

Investment committee decks should be the zenith of accuracy and clarity in decision-making material. They require layering AI efficiencies with governance rigor to prevent hallucinations from creeping into briefing materials.

1. Establish Clear Use-Cases for AI Outputs

Define exactly where AI-generated content is permissible:



- **Preliminary Research Basis:** Use AI for idea exploration, early screening, and summary drafting.
- **Not for Final Facts:** Never rely on AI as the sole source for deal specifics, valuations, or legal terms.

This ensures the team treats AI content as input, not gospel.

2. Layer In a Robust Fact Checking Process

Develop a systematic approach to validate AI-generated content including:

1. **Source Cross-Verification:** Match AI-derived facts against primary documents—financial statements, legal agreements, LPA excerpts, and validated data providers.
2. **Manual Review of Summaries:** Experts dissect AI summaries to confirm accuracy and context, especially for nuanced private market sectors.

3. **Track Changes and Source Notes:** Maintain transparent records of all modifications and source attributions, enabling downstream auditors and reviewers to trace origins.

3. Empower the Final Reviewer Role

Assign a senior investment professional—ideally with deep sector or private markets knowledge—to be the final gatekeeper for IC decks. Their responsibilities include:

- Validating key findings and assumptions
- Ensuring AI content aligns with known facts and strategic frameworks
- Rejecting or flagging any ambiguous or unverifiable AI statements
- Confirming that all AI summaries include explicit disclaimers about source and proof status

Essentially, this final reviewer acts as the human firewall against hallucination contamination.

Practical Workflow: From AI Draft to Audit-Ready IC Deck

Step Activity Tools/Controls Outcome
1 Preliminary Research with ChatGPT / Gemini Defined prompts, logging AI versions used, preliminary info only Initial ideas and summaries for deeper examination
2 Manual Extraction & Cross-Checking Primary documents, trusted databases (e.g., PitchBook, CapIQ), legal counsel review Verified facts and metrics underpinning analysis
3 Draft IC Deck Creation AI-assisted summarization with clear source attribution Consolidated deck with transparent reference trail
4 Final Reviewer Scrutiny Senior investment lead with checklist & red flags for hallucinations Clean, reliable, audit-ready IC deck
5 Committee Presentation & Documentation Version control, archival of reviewed versions, notes on AI usage Defensible record of analysis and decision rationale

Additional Tips: Avoiding AI Hallucinations in Family Office Settings

- **Keep a “Things AI Got Wrong” Log:** Document errors and inconsistencies encountered with AI tools to refine prompts and internal guidelines.
- **Don’t Trust AI Screenshots as Proof:** Screenshots can be fabricated or reflect earlier errors; always verify facts before inclusion.
- **Beware Vague Claims About AI Models:** Avoid accepting statements like “using latest GPT4+” without understanding the model’s training context and limitations.
- **Train Investment Teams:** Educate analysts and associates on AI’s strengths and pitfalls so they approach outputs critically.

The Bottom Line: AI Can Help—but Human Control Must Prevail

Think about it: the integration of AI tools such as ChatGPT and Google Gemini within family office IC processes is a leap [grok AI mistakes](#) forward for efficiency and insight generation. Yet, the seductive persuasiveness of AI hallucinations, especially in the opaque private markets, makes unguarded reliance dangerous. By instituting rigorous IC deck controls, a multiphase fact checking process, and entrusting a knowledgeable final reviewer to confirm accuracy, family offices can harness AI’s power while safeguarding their fiduciary integrity and client trust.

Remember: before adopting any AI tool or output into final investment decks, always ask yourself, *“What would we show an auditor?”* If confident that every claim can be traced, validated, and defended, the deck is ready. If not, it’s back to diligent human verification—a principle every UHNW family office should embrace.