

In the evolving landscape of wealth management and private markets investing, artificial intelligence (AI) tools like ChatGPT and Google Gemini have rapidly become go-to solutions for due diligence and preliminary research. Yet, as enticing as these tools are, especially for high-net-worth families and their advisors, their deployment in private fund research warrants a careful, nuanced discussion.

This article explores the benefits, risks, and practical considerations around allowing clients—particularly ultra-high-net-worth (UHNW) and family office investors—to use AI for private fund research. From the nature of AI hallucinations to the severe scarcity of verifiable private market data, the key themes uncovered here are critical for advisors focused on client education and responsible technology governance.

Understanding AI Capabilities in Private Fund Research

ChatGPT by OpenAI and Google Gemini represent two of the most advanced large language model (LLM)–based AI tools currently available. Both use extensive [trust but verify AI](#) training on diverse datasets to generate natural language text reminiscent of human analysis.



Use cases in family office contexts often include:

- Conducting preliminary research on fund managers or strategies
- Parsing complex private placement memoranda (PPMs), subscription agreements, and other lengthy documents
- Generating summaries that distill dense legal and financial content into digestible insights
- Assisting in the organization and classification of due diligence materials

These applications can speed up repetitive tasks, easing cognitive load and allowing investment teams to focus on higher-value analysis. Clients who want to independently understand private funds may find AI an appealing shortcut.

The Problem of AI Hallucinations: Persuasive but Risky

Nonetheless, the so-called “hallucination” phenomenon inherent to LLMs poses significant risk. Hallucinations occur when AI confidently generates plausible-sounding information that is incorrect, fabricated, or misleading.

Why Hallucinations Are Especially Dangerous in Private Fund Research

- **Private Markets Data Scarcity:** Unlike public equities where vast amounts of verified data exist, private funds often have limited accessible information. This scarcity makes it easier for AI to “fill in gaps” with fabricated details.
- **Complex Fund Structures:** Private funds can have complex fee arrangements, side letters, and bespoke governance that generic AI models are unlikely to fully capture or interpret correctly.
- **Persuasive Language:** AI outputs are linguistically fluent, making hallucinations difficult for non-experts to detect—clients may take the output at face value without critical scrutiny.

For example, a client who asks ChatGPT about a private equity fund’s performance history may receive a smooth narrative including fictitious track records or manager commentary. Even an advisor reviewing AI-generated summaries must maintain vigilance to identify such errors.

Faithfulness vs. Factual Hallucinations: Why It Matters

In AI terminology, *faithfulness* refers to how consistent a generated summary or explanation is with the source data or documents. A *factual hallucination* is when AI includes information not supported by the input.

In private fund due diligence, faithfulness guarantees are paramount. If an AI tool generates a “summary” of a private debt fund’s offering documents but invents terms or misstates risks, it undermines trust and may expose clients and advisors to compliance problems.

Governance around AI use must emphasize measuring and communicating faithfulness metrics, ideally empowering users with disclaimers about possible hallucinations and guidance on verification.

Family Office Use Cases: How AI Can Add Value

Despite risks, AI tools can be valuable when employed thoughtfully within family offices’ research workflows. Examples include:

1. **Preliminary Research:** AI can quickly compile background details on managers, investment strategies, or market themes from unstructured data sources.
2. **Document Parsing and Summarization:** Parsing PPMs, fund fact sheets, and subscription agreements to generate first-level summaries that highlight key fee terms, lockup periods, or investment restrictions.
3. **Data Organization:** Classifying and tagging due diligence documents for easier retrieval and review by internal teams.

Importantly, in all these use cases AI outputs should be regarded as starting points, not final conclusions. Human review remains essential to validate, challenge, and contextualize findings.

Client Education and Hallucination Warnings

Allowing clients direct access to AI tools is a growing trend, but it calls for robust education and oversight. Family offices and multi-family office providers considering this must:

- **Implement Clear Hallucination Warnings:** Before clients use ChatGPT or Google Gemini for private fund research, provide explicit disclaimers that AI outputs may contain errors, omissions, or invented details.
- **Train Clients on Critical Thinking:** Educate users to treat AI-generated content as preliminary, requiring confirmation against official fund documents and manager communications.
- **Set Guardrails Around Data Sensitivity:** Discourage inputting confidential or non-public information into public AI platforms to protect privacy and ensure compliance.
- **Offer Human-in-the-Loop Services:** Provide advisory support where staff review and validate AI-generated research to safeguard accuracy before client consumption.

Comparing ChatGPT and Google Gemini for Private Fund Research

Feature ChatGPT Google Gemini Training Data Scope Broad but with known knowledge cutoff, relies on static datasets Integrates Google's real-time data & search, potentially more up-to-date Integration with Private Data Requires manual uploading or API linking Designed to better integrate with Google Cloud and enterprise datasets Hallucination Risk Moderate to high, especially on niche private market topics Claims improved factuality, but independent validation still crucial Summary & Parsing Abilities Strong natural language summarization Advanced contextual understanding, leveraging Google's search algorithms Client Accessibility Widely available, straightforward UI More limited access, positioned for enterprise adoption

While Google Gemini promises reductions in hallucination rates through integration with up-to-date data and advanced retrieval methods, neither tool can fully solve the core issue of private market opacity and data sparsity.

Practical Recommendations: What Would We Show an Auditor?

Before implementing AI solutions for client use in private fund research, family office risk leads must ask fundamental questions, including:

- Can we demonstrate the provenance and review process for AI-generated research outputs?
- Are we logging AI interactions to allow audit trails and post-hoc analysis?
- Do we have policies restricting the types of data shared with AI tools, especially commercially sensitive or material non-public information?
- Can we provide standardized hallucination warnings alongside any AI output delivered to clients?
- Is there a defined workflow that ensures final investment decisions rely on validated inputs, not just AI summaries?

Showing auditors a transparent governance framework around AI use is essential to manage operational, reputational, and compliance risks.

Conclusion

Allowing clients to use AI tools like ChatGPT or Google Gemini for private fund research is a double-edged sword. The speed and convenience AI offers can democratize access and foster engagement, yet hallucination risks and private market data scarcity demand a cautious, well-controlled approach.

Key takeaways for family offices and advisors:



1. Use AI for preliminary research and document summarization—but always pair outputs with human expert review.
2. Educate clients rigorously on hallucination risks and why AI-generated content is not definitive investment advice.
3. Implement clear hallucination warnings and maintain audit logs of AI-assisted client interactions.
4. Review AI deployment strategies through a compliance and risk lens, showing auditors how outputs are vetted.

In the current state of AI and private market data, trust—but verify—should be the guiding motto.

This article reflects the author's 12 years of experience supporting multi-family offices in Canada, emphasizing rigorous data governance and thoughtful adoption of AI technologies.