

As family offices increasingly explore artificial intelligence (AI) tools like ChatGPT and Google Gemini for generating research notes, compliance teams face new challenges balancing the benefits of speed and efficiency against the pitfalls of accuracy and traceability. AI-generated research offers a promising avenue for preliminary research, parsing voluminous reports, and distilling complex content into digestible summaries — but the risks of hallucinations, especially in the context of scarce private markets data, demand rigorous oversight. Here's a story that illustrates this perfectly: learned this lesson the hard way..



In this article, I outline what compliance teams should prioritize when reviewing AI-generated research notes to mitigate risks and maintain an effective **audit trail** and **source documentation**.

Understanding the Risk Landscape: AI Hallucinations Are Persuasive and Risky

One of the foremost risks with tools like ChatGPT and Google Gemini is their tendency to generate plausible but false or misleading information, known in the industry as “hallucinations.” These hallucinations are particularly dangerous because the language models produce confident, fluent prose that can easily mislead an unwary reader. Ignoring this risk is perilous, especially for family offices making investment decisions on behalf of ultra-high-net-worth individuals.

Compliance must internalize these realities:

- **Hallucinations are not bugs but inherent limitations:** Large language models generate outputs based on patterns in training data; they do not “know” facts or verify information.
- **Persuasive, but not always truthful:** AI will generate answers that *sound* correct, even if they contain errors or fabricated sources.
- **Private market data scarcity compounds risk:** Since private markets data is limited, less factual information exists for AI models to train on, increasing hallucination frequency in this sector.

Family Office Use Cases: Where AI-Generated Research Notes Are Leveraged

Understanding typical family office use cases for AI-generated notes helps frame compliance priorities. Common applications include:

1. **Preliminary Research:** Quickly gathering background information about an industry, company, or strategy before diving into deeper due diligence.
2. **Parsing Lengthy or Complex Reports:** Condensing dense financial statements, regulatory filings, and market commentary into executive summaries.
3. **Generating Summaries:** Creating concise, readable overviews of various data sources or discussion points for stakeholders.

While AI tools increase efficiency and reduce manual workload, integrating their output into formal documentation requires ensuring outputs are factually accurate and verifiable.

Faithfulness vs Factuality: Why They Matter

Two related but distinct concepts are critical when evaluating AI outputs for compliance:

Term	Definition	Compliance Implication
Faithfulness	The degree to which the AI's output accurately reflects the input data or source material. Ensures summaries and notes do not distort or omit key facts when condensing source documents.	
Factuality	The objective truthfulness of the information presented in the AI-generated content. Requires verification against authoritative sources since hallucinations may produce false facts.	

Faithfulness evaluates whether the AI is correctly paraphrasing or summarizing known inputs, while factuality asks whether the contained facts themselves are true. Compliance teams must demand both for reliable research notes.



Key Compliance Review Elements for AI-Generated Research Notes

To manage risks while benefiting from AI, compliance teams should adopt a structured review process with attention to the following areas:

1. Establish a Robust Audit Trail

An effective audit trail is fundamental. It should record:

- The specific AI tool used (e.g., ChatGPT version, Google Gemini iteration).
- Timestamped queries/prompts submitted.
- Full AI-generated outputs stored in perpetuity for later review.
- Subsequent human-reviewed edits and annotations.

This trail enables reconstructing how conclusions were reached and defends against regulatory or internal queries.

2. Insist on Source Documentation and Reference Transparency

AI's easiest failure point is unreferenced or fabricated citations. Compliance should require:

- Clear pointers to original human-verified sources for key data points or claims in research notes.
- Verification protocols ensuring cited sources exist and contain supporting facts.
- Prohibition on presenting AI-generated "facts" without linking to authentic documentation or authorized databases.

3. Use AI as a Support Tool, Not a Primary Source

Policies should <https://canadianfamilyoffices.com/special-reports/technology-sr/ai-has-the-makings-of-a-powerful-tool-for-family-offices-but-frequent-and-highly-convincing-hallucinations-are-proving-troublesome/> clearly state AI-generated notes are *starting points* for human analysts, not substitutes for human diligence. This mindset helps prevent overreliance on AI outputs. Analysts must:

- Cross-check AI summaries against original reports and data.
- Flag inconsistencies or suspicious statements for deeper investigation.
- Supplement AI notes with expert commentary and context.

4. Train Staff on AI Limitations and Error Types

Intelligent use of AI requires understanding its limitations. Compliance and investment teams must be trained on:

- Understanding different types of hallucinations.
- Recognizing red flags like overly confident wording or missing citations.
- Knowing when to escalate unclear or possibly incorrect outputs.

5. Implement Continuous Monitoring and Feedback Loops

Compliance should establish ongoing sample audits of AI-generated research notes to detect error trends or suspicious outputs early. Corrective feedback loops with users and vendors enable:

- Identifying persistent hallucination patterns in specific knowledge domains (e.g., private markets).
- Updating internal best practices and prompt designs to enhance output quality.
- Ensuring vendor transparency on model improvements and data governance.

Special Considerations for Private Markets Data Scarcity

Family offices often operate heavily in private equity, real estate, and direct investments where public data is sparse. This scarcity raises the hallucination risk for AI-generated research dramatically because:

- AI models have limited factual training material in these areas.

- Consequently, hallucinations tend to be more frequent and impactful.
- Sources cited may be non-existent or outdated.
- Nuances unique to private asset classes may be misunderstood or oversimplified.

Compliance takeaway: Private markets research notes generated by AI require extra scrutiny, stringent source verification, and mandatory human validation before any decision-making use.

Summary: Best Practices Checklist for Compliance Review

Review Area Action Items Audit Trail Record tool versions, prompts, AI outputs, and human edits; store securely. Source Documentation Ensure verifiable citations; cross-check references against original documents. Human Oversight Use AI output as a starting point only; analysts validate and annotate. Staff Training Educate teams on AI capabilities, hallucination types, and error detection. Ongoing Monitoring Conduct periodic audits; feed findings back into AI use policies and vendor dialogues. Private Markets Focus Apply heightened validation for private asset research due to data scarcity risks.

Concluding Thoughts

AI tools like ChatGPT and Google Gemini are transforming how family offices perform research. They enable rapid synthesis and initial analysis that can free human capacity for higher-value tasks. Yet, the compliance function must remain vigilant in demanding transparency, accuracy, and provenance in AI-generated research notes. The seductive eloquence of AI hallucinations can mask serious errors—with costly consequences if left unchecked.

By enforcing strong audit trails, demanding verifiable sources, using AI as an assistive tool rather than a final arbiter, and applying extra caution in private markets, compliance teams can responsibly leverage AI's promise while preserving trust and protecting family office interests.