

In today's rapidly evolving technology landscape, artificial intelligence (AI) tools such as **ChatGPT** and **Google Gemini** have become go-to assistants for many professionals, including family office teams managing complex portfolios and private market investments. These tools promise fast, comprehensive answers and summaries, but what happens when their output conflicts with your own internal research? This is an increasingly common challenge that requires a thoughtful approach to conflict resolution, especially when managing data scarcity and the high stakes of ultra-high-net-worth (UHNW) investment decisions.

Understanding the Risk of AI Hallucinations

One of the foremost risks when relying on AI-generated content is the phenomenon known as *AI hallucination* — when the AI confidently produces inaccurate or fabricated information. These hallucinations are often *persuasive*, filled with seemingly authoritative language, making them particularly dangerous to uncritical readers.

For family offices, whose decisions impact multi-million dollar portfolios, relying on unverified AI output can lead to costly mistakes or regulatory complications. Unlike traditional databases or expert reports, large language models (LLMs) like ChatGPT and Google Gemini do not inherently fact-check or cite primary documentation unless specifically instructed and supported by governance procedures.

Why AI Hallucinations Occur More Frequently in Private Markets

Data scarcity is a major factor increasing AI output error risk in private markets. Unlike public equities or fixed-income assets, private equity, venture capital, and other alternative investments often have limited, inconsistent, or non-public data sources. AI models trained primarily on web content or publicly available datasets may struggle to provide quality output here.

As a result, these models may generate plausible but incorrect figures, dates, or relationships due to the lack of concrete data — a challenge family office researchers cannot overlook when they see conflicting outputs compared to their own due diligence or proprietary data warehouses.

Family Office Use Cases for AI Tools: Preliminary Research and Report Parsing

Despite the risks, AI tools hold undeniable value for family offices when used appropriately:

- **Preliminary research:** Quickly exploring unfamiliar sectors or companies to build a foundational understanding.
- **Parsing and summarizing reports:** Accelerating the digestion of lengthy analyst reports, fund documents, or regulatory filings.
- **Data synthesis:** Creating first-draft summaries for further internal review.

However, AI output is best viewed as an input into a *structured research workflow*, not a final source. Integrating tools like ChatGPT and Google Gemini can save time but must be combined with rigorous validation practices.

Faithfulness vs. Factual Accuracy: What to Keep in Mind

One useful distinction when evaluating AI output is between *faithfulness* and factual accuracy. Faithfulness refers to how consistently an AI's output reflects the input data it was provided, whereas factual accuracy concerns

whether that information aligns with reality.

Both ChatGPT and Google Gemini have mechanisms to improve faithfulness — they attempt to avoid outright contradictions of their given context — but neither fully guarantees factual correctness. This gap means that even text consistent with a prompt can still convey incorrect facts.

For family office research, this emphasizes the necessity of:

- **Cross-checking AI output against primary documents:** Fund prospectuses, audited financial statements, or signed investment memoranda.
- **Maintaining a clear source hierarchy:** Ranking internally vetted research above secondary summaries and AI-generated content.
- **Documenting discrepancies:** Keeping track of instances where AI conflicts with internal data, noting potential causes and resolutions.

Conflict Resolution Steps When AI Output Contradicts Internal Research

When faced with contradictory information, a systematic approach reduces risk and builds team confidence in AI-assisted workflows. The following steps are recommended:

1. **Pause and assess the nature of the contradiction:** Determine if the AI output conflicts with a key fact, an interpretation, or a summary nuance.
2. **Refer back to primary documents immediately:** For example, check official financial statements, regulatory filings, or signed agreements — the most authoritative sources your office holds.
3. **Check your internal source hierarchy:** Establish upfront that primary documents and in-house vetted research override AI summaries. AI should never replace core internal data.
4. **Use AI for hypotheses, not conclusions:** If AI suggests contradictory facts, treat those as prompts for deeper dives or queries to fund managers.
5. **Escalate complex conflicts:** Discuss discrepancies with portfolio managers, external counsel, or data vendors as appropriate.
6. **Document and log conflicts:** Maintain a living record of “things AI got wrong” (a personal practice I strongly recommend) to identify common hallucination triggers and avoid repeated errors.
7. **Evaluate the AI tool’s update and governance status:** Ask vendors for transparency about model versions, data recency, and validation protocols before relying on outputs.

Establishing a Clear *Source Hierarchy* Is Key

Source hierarchy means your research framework clearly ranks information by reliability, establishing a mental model for every team member. Here’s an example:

Rank Source Type Example(s)
1 Primary Documents Signed agreements, audited financial statements, regulatory filings
2 Internal Research & Due Diligence Original models, confidential manager diligence notes, proprietary databases
3 Authoritative Third-Party Reports Analyst reports, fund marketing materials from trusted sources
4 AI-Generated Content Summaries, hypotheses, exploratory research via ChatGPT or Google Gemini

Want to know something interesting? by institutionalizing this hierarchy, you can avoid the common trap of elevating an ai-generated summary over validated internal findings. This is especially critical with private markets where external data scarcity amplifies risk.

What Would We Show an Auditor? - A Governance Lens

A question I always ask myself before adopting any AI tool or trusting its output is: **“What would we show an auditor?”** Auditors require clear evidence trails, source documentation, and validation frameworks.

AI outputs alone are rarely suitable for audit or compliance purposes because they often lack citations or provenance. Any AI-generated insight must be backed by primary documents and logged into your approved internal data ecosystem. When you can’t demonstrate that linkage, the output is just a suggestive tool, not a final answer.



Final Thoughts: AI As a Helpful Collaborator, Not a Replacement

The emergence of LLMs such as ChatGPT and Google Gemini presents exciting opportunities to streamline family office research and reporting workflows. However, the persuasive nature of AI hallucinations and inherent challenges in private markets data require heightened vigilance.

By establishing clear conflict resolution processes, grounding your research in a defined source hierarchy, and cross-checking AI output against primary documents, your office can safely harness AI’s power without compromising accuracy or compliance.



Remember: AI should augment your expertise, not replace it. When in doubt, trust your internal rigor and the ultimate question: *"What can we confidently show an [human in the loop AI](#) auditor?"*

Summary Checklist for Conflict Resolution in AI vs. Internal Research

- Identify if AI output contradicts critical facts or interpretations
- Confirm disagreement using primary source documents immediately
- Recall the agreed-upon source hierarchy: primary docs > internal research > AI output
- Use AI-generated info as prompts, not conclusions
- Consult experts or external partners on unresolved conflicts
- Document discrepancies to improve future governance
- Evaluate and question AI vendor transparency and model governance
- Consider audit-readiness before applying AI outputs in official reporting