Frontiers in Business, Economics and Management
ISSN: 2766-824X | Vol. 23, No. 2, 2026

The Augmented Advisor Model in Wealth Management: A Qualitative Exploration of GenAI Application Efficiency and Human Trust Boundaries

Ronghe Liu *

School of Letter and Science, University of California, Davis, California 95618, USA

* **Corresponding author**: (Email: liuronghe0901@outlook.com)

Abstract: In the future, the Wealth Management sector may have an upgraded form known as "enhanced advisors". It is expected that generative AI will assist people's work with reasoning abilities; however, they will not fully replace human la bour. This paper qualitatively explores the new form through two inter-related aspects: application efficiency and the limits of people's trust. Based on relevant academic literature and institutional reports published from 2019 to 2026, this paper examines how l arge language models (LLMs) and agentic AI can improve operational efficiency in three areas: back-office compliance, front-end personalisation and client communication workflows; then analyzes unevenly distributed effects using the "prompt dividend" mechanism and supply-side model bias. On this basis, a triangular advisor-client-AI trust mechanism is used to model the generation and maturation process of trust from "avoidance" to "acceptance". Through analysis, although the number of GenAIs has increased through certain means, building trust requires explanation, human management, and organizational supervision. This paper proposes that a zero-trust governance architecture should be integrated with a human-in-the-loop mechanism to address the disconnect between efficiency potential and trust realisation.

Keywords: Generative Artificial Intelligence; Augmented Advisor; Wealth Management; Human–AI Trust; Human-in-the-Loop.

1. Introduction

In recent years, for a long time it has been restricted in terms of three factors, namely too much attention to professionals; tight regulation; strong inter-dealer relationships [1]. In recent years, algorithmic robo-advisors have entered the market through automation of investment selection and distribution platforms for ordinary individuals in finance [2]. However, these earlier systems were based on static and rule-driven models without being able to handle unstructured data or conduct detailed context-based conversations with users.

The emergence of generative artificial intelligence, driven by transformer-based large language models (LLMs) in recent years, has entirely reversed the course of change. Since the last generation, it cannot only comprehend language but also perform logical reasoning, produce content, and process multimodal information, thus opening up a brand-new direction for finance [3]. With the development of technology, there has emerged a form known as "augmented advisor": artificial intelligence will assist people rather than completely substitute for them; moreover, it can help with tasks such as analyzing data and facilitating communication between humans [4].

Therefore, this type of guarantee is realizable. Industry forecasts indicate that Generative AI can bring up to hundreds of billions in annual added value for the banking industry by improving productivity [5]. Based on field investigations, additional human supervision with an AI-based advisory system enhances clients' trust in investment advice at critical junctures [6]. Nevertheless, this efficiency potential has yet to address the trust gap. Based on investigations, although some clients have expressed satisfaction with the efficient response speed of AI for simple questions, they are still skeptical of autonomously produced comprehensive financial planning suggestions [7].

The conflict between speed and confidence is the core issue of this paper. How does the augmented advisor model alter the efficiency pattern of wealth management, and what are the psychological and institutional limits of human belief in a human-AI cooperative framework? Most existing research on robo-advisors has shown that algorithm-based portfolio construction can enhance diversification and reduce behavioral biases in retail investors [2]. However, this conclusion is based on earlier research conditions. Meanwhile, the emerging literature on LLMs in finance focuses mainly on technical indicators [3, 8], with relatively little attention paid to human-related factors such as trust construction and institutional governance at present. This paper attempts to connect the two.

The remainder of this paper is organized as follows. Section 2 reviews the relevant literature. Section 3 proposes a conceptual framework linking efficiency and trust. Section 4 presents the analysis and discussion. Section 5 offers implications and recommendations. Section 6 concludes.

2. Literature Review

2.1. GenAI in Financial Services: Technological Foundations

The technical backbone of the augmented advisor is built around transformer networks and their self-attentive mechanism; it can find long-range semantic relationships in complicated financial documents, including earnings announcements, regulatory documents, and research reports [3]. Domain-specific LLMs have achieved some results through targeted pre-training. BloombergGPT, which has been pre-trained in a selected financial domain, demonstrated higher accuracy than general-purpose models of comparable size for financial sentiment analysis, named-entity recognition, and question-answering tasks [3]. Broader surveys of LLM applications in finance have listed use cases such as investment analysis summarisation, regulation text interpretation, client communication writing, risk factor identification, and other tasks [8].

Recently, the development towards agentic AI (systems with autonomous planning ability, continuous state monitoring capabilities, and immediate tool invocation functions) has extended the application range of AI from reaction mode to proactive operation. The agency system can automatically perform account monitoring to be notified of rebalancing if there has been a variation during human life changes; lastly, give specific course-of-action recommendations [9]. Based on industry analysis, most wealth management companies have recognised the potential of agentic AI [4].

2.2. Efficiency Evidence in Wealth Management

Efficiency improvements brought about by GenAI in wealth management can be divided along the value chain. In back-office operations, the application of AI-driven auditing systems has changed from sample verification to full population evaluation, which enables real-time anomaly detection and significantly reduces false alarms in anti-money laundering monitoring [5, 10]. Front-end interaction in terms of personalized service is a kind of "scalable individualisation", which could not be provided for ultra-high net worth individuals before [4]. However, recent research has shown that the efficiency gains are not evenly distributed; in terms of the effectiveness of AI-generated advice, the difference in users' level of financial literacy and prompting ability also needs to be considered, so that some people will obtain more wealth growth dividends [11].

2.3. Trust in human-AI decision-making

Studies have long examined trust in automated systems. Based on [12] by Lee and See, trust can be regarded as an agent's intent to assist in achieving its own goals under circumstances of uncertainty and vulnerability. Two opposing states occur among users regarding algorithmic recommendation systems. Algorithm avoidance, that is, rejecting an algorithm-generated suggestion upon discovering only one error, was first identified by Dietvorst et al. [13] but then validated in delegated investments [14]. Consequently, the generation of an algorithm appreciation tendency tends to occur under conditions where decision complexity exceeds individuals' cognitive limits [15].

Field experimental evidence has provided particularly compelling insights. Yang et al. [6] conducted a large-scale field experiment on human–AI collaboration in investment advisory and found that human presence alongside AI significantly increased clients' adoption of investment recommendations. Crucially, this effect appeared to stem less from perceived quality improvements, and more from the emotional and accountability signals conveyed by having a human "in the loop." Recent integrative reviews suggest that trust in AI is a multidimensional phenomenon shaped by perceived competence, explainability, accountability, privacy, and fairness [7].

2.4. Research Gap and Positioning

Despite growing scholarly attention, the intersection of "wealth management + GenAI + augmented advisor" remains underexplored in peer-reviewed literature. Most of the existing evidence needs to be moved across neighboring fields, including hybrid robo-advisory and general human-AI collaboration, and set reasonable boundaries. This paper synthesises efficiency evidence and trust mechanisms within a general analysis framework for the enhanced advisory environment.

3. Conceptual Framework

This paper proposes a two-dimensional framework, the Efficiency-Trust Framework (ETF), to analyse the augmented advisor model, as shown in Fig. 1.

This paper splits the efficiency part into supply-side efficiency (intrinsic ability enhancement and bias correction for AI model output) and demand-side efficiency (how users' financial knowledge level and reminding ability impact the quality of AI-generated advice). The decomposition is based on empirical evidence that users' factors contribute to about 67% of the variation in advice quality; meanwhile, model-specific biases account for the remainder [11].

Trust can be examined through a combination of three parties: the human advisor, the client and the intelligent system. Unlike traditional dyadic trust models, the three-dimensional structure is capable of depicting that the critical mediator of trust in this system is an intermediary who has both sides' recognition. In addition, it also shows that there are both trustee-creditor and AI-creditor relationships [6]. The framework suggests that the realisation of the efficiency potential of GenAI falls within the trust boundary; once its gain surpasses this threshold, it will not be trusted directly but needs institutional governance measures to obtain approval.

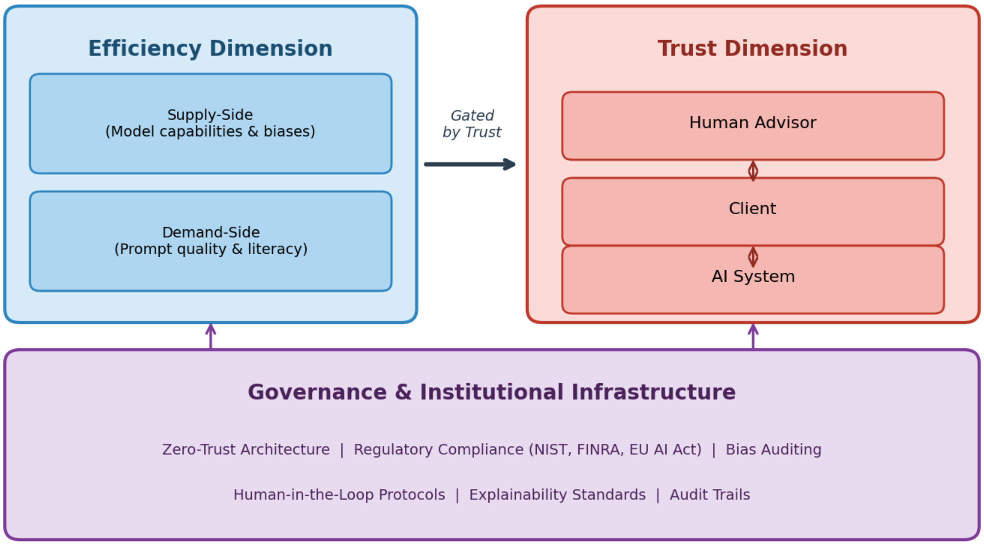

Figure 1. The Efficiency–Trust Framework for the Augmented Advisor model

4. Analysis and Discussion

4.1. Efficiency Gains Across the Value Chain

Evidence of the effectiveness contributions of GenAI covers both backend and frontend processes. Table 1 presents the main inefficiency effects determined from previous research.

Table 1. Summary of GenAI Efficiency Impacts Across the Wealth Management Value Chain

Function	Traditional Bottleneck	GenAI Enhancement	Source
Financial Audit	Sample-based, time-intensive	Full-population real-time analysis	[5, 10]
AML/Fraud Detection	Fixed rules, high false-positive rates	Context-aware anomaly detection	[5, 10]
Client Communication	Manual drafting, limited personalization	Automated, context-sensitive narratives	[4, 6]
Research Synthesis	Hours of manual reading	Multi-document summarization	[3, 8]

Efficiency gains are most evident for text-heavy retrieval-type tasks; that is, the areas with an explicit superiority of transformer-based models over LLMs. Tasks that demand value judgment, emotion recognition or strategy innovation remain within the scope of human strengths relative to AI advisory teams so far.

4.2. The Uneven Distribution of Efficiency: Prompt Dividends and Supply-Side Bias

Recent studies show that the benefits of a GenAI system's efficiency are unevenly distributed among users. On the demand-side, those who are more financially literate have created more orderly, restrictive prompts and thus elicited better-aligned responses from artificial intelligence with existing normative frameworks of financial planning. Based on simulations of life-cycle models, there are generally 4%-5% retirement wealth gap issues after an average 40-year-long process under the "prompt dividends" mechanism [11].

There is an absolute error in the model's output for identically presented sets of financial data. Research indicates that LLMs have a tendency to make more cautious investment recommendations when processing profiles that are identified as belonging to female investors, and such biases contribute roughly three-quarters of the variance in AI-generated opinions [11]. Table 2 summarizes these factors.

Table 2. Demand-Side and Supply-Side Factors Affecting AI Advisory Outcomes

Factor Type	Specific Factor	Mechanism	Estimated Impact
Demand	Financial literacy	Determines prompt sophistication and constraint specification	4–5% retirement wealth gap [11]
Demand	AI usage experience	Enables iterative correction of model hallucinations	Experienced users accumulate more savings [11]
Supply	Demographic bias	Model weights produce differential risk anchoring	~1/3 of total advice variation [11]

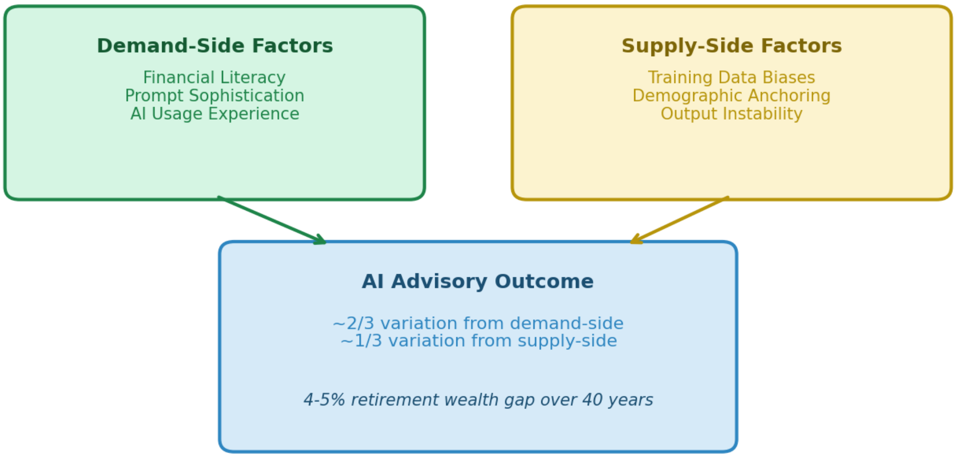

Figure 2. Interaction model of demand-side and supply-side factors in AI advisory outcomes

4.3. Trust Dynamics in the Triadic Relationship

The trust dimension in the Augmented Advisor model is not explained by traditional dyadic trust models. Introducing AI into the advisory relationship generates a tripartite dynamic of Advisor-Client-AI where each party holds different trust calculus.

4.3.1. The Three-Stage Trust Evolution

Based on field evidence of bank institutions' application of AI technology, the process of people's confidence in artificial intelligence is mainly divided into three stages. In the first stage (vulnerability avoidance), advisors identify AI technology as an external threat to their profession and clients, thus taking physical and functional isolation measures for AI tools in non-core back-office positions. In the second phase (vulnerability safeguarding), as confidence gradually increases in the predictive capabilities of artificial intelligence, it leads to higher frequency of use among customers for handling operations based on AI predictions; however, this requires dual-signoff technology that mandates approval from both parties individually first. In the third stage (vulnerability acceptance), advisors acknowledge that AI's ability to process heterogeneous data flows exceeds humans' capabilities, so advisors are regarded as "ultimate reviewers" of AI output results and "deep resonators" who interpret clients' emotional states, thus achieving a true state of human-AI symbiosis. As shown in Table 3.

Table 3. Three-Stage Evolution of Advisor Trust in AI Systems

Stage	Advisor Behavior	Psychological Driver	AI Usage Pattern
1. Vulnerability Avoidance	Defensive, restrictive	Replacement fear, opacity anxiety	AI confined to back-office
2. Vulnerability Safeguarding	Selective, conditional	Need for control, reputation protection	AI with mandatory human sign-off
3. Vulnerability Acceptance	Integrative, collaborative	Recognition of complementary strengths	Advisor as final reviewer

4.3.2. Moderating Variables

Trust levels with different groups of people vary. According to research in consumer psychology, "maximisers", who are always eager for optimisation, tend to appreciate algorithms more strongly than do "satisficers" [15]. Maximisers experience high cognitive load in the face of a large number of options on modern financial markets because of GenAI's aggregate function to alleviate it. Explainability also needs to serve as an intermediary between all participants in the three-part cooperative mechanism [7, 12]. The OECD recommended adopting a realisation process known as the "multilayer explanation" for financial organisations; that is, under different circumstances of disclosure at varying levels [5].

4.4. Governance as an institutional foundation for trust

As agentic AIs progress from suggesting to completing transactions, governance provides the core foundation of trust construction. NIST's AI risk management framework for generative artificial intelligence organises the types of risks into a certain order and classifies confounding, data privacy violations, harmful bias and information security problems as its primary concern [10]. According to this design, a zero-trust governance structure should be built for physical separation of the decision-making and execution units, referring to Figure 3. The design includes the following three parts. 1. Managed Data Layer, which ensures that inputs are correct; 2. Sandboxed Reasoning Layer for calculating a hallucination risk score and requiring certification from humans for final verification;

FINRA stated at present that supervisory and recording duties, fair-trade regulations, etc., still apply to newly developed securities businesses based on new technology enablers; companies must set up an AI monitoring system with the functions of real-time audit capabilities as well as immediate error-circulation capacity. The EU's AI Act and the OECD guidelines offer each other supplementary regulation bases to some extent [5].

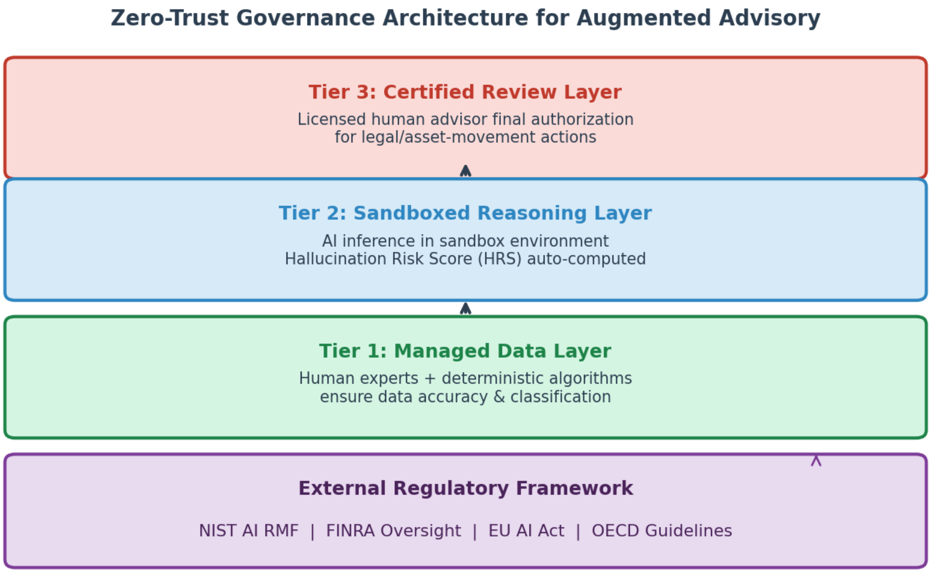

Figure 3. Three-tier zero-trust governance architecture for the Augmented Advisor model

5. Implications and Recommendations

5.1. For Practitioners

Following the three rules established by financial institutions employing the Augmented Advisor approach. First, deploy GenAI initially in text-intensive back-office functions where efficiency gains are most robust and trust requirements are lowest. Secondly, introduce an absolute requirement of human intervention for all clients facing AI output to achieve quality assurance at the level of human oversight in order to enhance acceptance by users. Thirdly, provide client financial education and increase the equity of quality prompts.

5.2. For Regulators

The regulatory system needs to adapt to these particular risks of the enhanced version. It should require the establishment of a complete audit trail of AI decision paths; set out explanation layer standards; develop an evaluation protocol for supply-side bias. There should be no regulatory gap created by AI; the original fiduciary duties applicable to human advisors still apply to those of AI [10].

5.3. Limitations and Future Research

This study has several limitations. As a qualitative exploration based on literature synthesis, it does not generate primary empirical data. Many cited studies originate from adjacent domains and their transferability requires validation. Future research needs to pursue cross-cultural studies on trust boundaries, longitudinal tracking of prompt dividends, and standardised benchmarks for assessing whether efficiency gains translate into measurably better client outcomes.

6. Conclusion

Augmented advisor is transformed into "Digital Wealth Management Partner", and then GenAI improves overall efficiency by more than 14 times for heavy text-based data preprocessing, as well as massive personalisation, with over twice the effectiveness in some areas. These gains are unevenly distributed across different situations due to varying degrees of effectiveness of demand-side prompts and biases in supply-side models.

In terms of trust building on the trust dimensions, the triangular relationship model shows that this type of trust is an evolving process shaped by professional identities, cognitive load, explainability, and institutional governance. Vulnerability avoidance, security protection, and eventually acceptance form the three-stage development of trust in Artificial Intelligence (AI).

Therefore, in conclusion, the future development direction of Augmented Advice must balance technology rationality with emotional resonance. It needs artificial intelligence to handle certain cognition-heavy tasks; at this time, human wisdom will take over and make a decision autonomously after all. To achieve this goal requires continuous investments in people, institutions and regulations to build trust in trustworthy human-AI collaboration.

References

[1] F. D'Acunto, N. Prabhala, and A. G. Rossi, "The promises and pitfalls of robo-advising," The Review of Financial Studies, vol. 32, no. 5, pp. 1983–2020, 2019. doi: 10.1093/rfs/hhz014

[2] A. G. Rossi and S. Utkus, "The diversification and welfare effects of robo-advising," Journal of Financial Economics, vol. 157, Article 103869, 2024. doi: 10.1016/j.jfineco.2024.103869

[3] S. Wu et al., "BloombergGPT: A large language model for finance," arXiv:2303.17564, 2023. URL: https://arxiv.org/abs/2303.17564

[4] A. Ramachandran, "Advanced AI in wealth management: Functional transformation and intelligent orchestration across the client lifecycle," ResearchGate Preprint, 2025. URL: https://www.researchgate.net/publication/390947298

[5] OECD, "Generative artificial intelligence in finance," OECD Artificial Intelligence Papers, No. 9, 2023. doi: 10.1787/ac7149cc-en

[6] C. (L.) Yang, K. Bauer, X. Li, and O. Hinz, "My advisor, her AI, and me: Evidence from a field experiment on human–AI collaboration and investment decisions," Management Science, vol. 72, no. 1, pp. 242–264, 2025. doi: 10.1287/mnsc.2022.03918

[7] S. Afroogh et al., "Trust in AI: Progress, challenges, and future directions," Humanities and Social Sciences Communications, vol. 11, Article 1071, 2024. doi: 10.1057/s41599-024-04044-8

[8] Y. Nie et al., "A survey of large language models for financial applications: Progress, prospects and challenges,"

arXiv:2406.11903, 2024. URL: https://arxiv.org/abs/2406.11903

[9] Y. Dong et al., "Large language model agents in finance: A survey bridging research, practice, and real-world deployment," in Findings of EMNLP 2025, pp. 17889–17907, 2025. doi: 10.18653/v1/2025.findings-emnlp.972

[10] NIST, Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile (NIST AI 600-1), 2024. doi: 10.6028/NIST.AI.600-1

[11] T. Choukhmane, T. de Silva, W. Lin, and M. Akuzawa, "How good is generative AI personal financial advice?" Working Paper, Nov. 2025. URL: https://www.timdesilva.me/files/papers/llm_advice.pdf

[12] J. D. Lee and K. A. See, "Trust in automation: Designing for appropriate reliance," Human Factors, vol. 46, no. 1, pp. 50–80, 2004. doi: 10.1518/hfes.46.1.50_30392

[13] B. J. Dietvorst, J. P. Simmons, and C. Massey, "Algorithm aversion: People erroneously avoid algorithms after seeing them err," Journal of Experimental Psychology: General, vol. 144, no. 1, pp. 114–126, 2015. doi: 10.1037/xge0000033

[14] M. Germann and C. Merkle, "Algorithm aversion in delegated investing," Journal of Business Economics, vol. 93, pp. 1691–1727, 2023. doi: 10.1007/s11573-022-01121-9

[15] J. M. Logg, J. A. Minson, and D. A. Moore, "Algorithm appreciation: People prefer algorithmic to human judgment," Organizational Behavior and Human Decision Processes, vol. 151, pp. 90–103, 2019. doi: 10.1016/j.obhdp.2018.12.005