Utility Difference Analysis of Single-Agent and Multi-Agent Large Language Models in Stock Investment Decision-Making

Yichuan Luo

College of Business, Shanghai University of Finance and Economics, Shanghai, 200433, China

2024210898@stu.sufe.edu.cn

Abstract. The area of financial investment decision-making is undergoing significant changes with the help of artificial intelligence and especially large language models (LLMs). The current state of research has largely been dedicated to the single agent application paradigm, this paper proposes and analyzes an innovative multi-agent LLM architecture. The research has developed a two-level structure of the division of labor into a specialist and a centralized decision-making site. There are three expert agents in the analysis layer: Value, Growth, and Macro. One of the agents of this multi-agent system happens to be the Chief Investment Officer (CIO) agent who oversees dynamic integration and the results of the 24-year backtest of a Nasdaq 100 index fund prove that the multi-agent system is much superior to the single-agent system on all the key performance indicators. It has grown by 111.8 percent in total return to 1482.31, its Sharpe ratio has almost doubled, the biggest drawdown has diminished by 36.9 percent, and it has effectively turned a negative value of Alpha to a positive value. This suggests that the counterbalancing of cognitive bias with the help of the multi-agent architecture will prove more efficient in generating excess returns and managing the risk.

Keywords: Large Language Models (LLMs), Multi-Agent Systems, Stock Investment Decision-Making, Quantitative Investment, Artificial Intelligence in Finance.

1. Introduction

Over the past years, financial investment sector has been radically transformed by the artificial intelligence (AI), particularly the large language models (LLMs), especially the information-intensive, decision-complex stock market (Cao et al., 2025; Ding et al., 2024). Their superior natural language understanding and generation abilities allow extracting important information and market moods out of unstructured input volumes of data such as news, financial reports, and social media in ways previously unimaginable (news, financial reports, and social media) to offer an unprecedented depth of analysis to quantitative approaches to investment while providing new dimensions quantitative strategies (Kelly, 2023). Nevertheless, most of the available literature has dealt with a single-agent paradigm, in which a single LLM model is used as a one-stop analyst. Although the potential of this technique has already been sufficiently proved, its limitations are also becoming apparent as well. On the one side, the performance of one model can ensure a decline in investment horizons in the long term and a large market sphere (Li et al., 2025). Alternatively, the cognitive biases, including overconfidence and anchoring effects, can be held (and even magnified) within the very LLMs, which will result in weakness of decisions (Malberg et al., 2024).

These are the limitations that are a significant weakness in modern research: how can the system of AI investment be created that will be more robust and adaptive and outperform the strength of an individual agent? The single-agent model is adversely affected especially by the complex and ever-evolving money markets that have a shortcoming of a mono-dimensional information processing mindset and lack of internal controls and balances. According to theoretical studies, a diversification and hedging of cognitive biases of separate agents can be achieved by introducing multi-agent systems with specialization of roles and collective intelligence (Grosu, 2025). This idea is consistent with one of the classical theories in finance, which is Agent-Based Computational Economics (ACE), which focuses on the simulation of interactions of heterogeneous agents to study the dynamics of a complex market (Wall and Leitner, 2021).

It is on this basis that this paper proposes and justifies a new, multi agent system based on the LLM used to invest in stocks. The system even mimics the collaborative pattern of a professional investment company, with its central design being that of "specialized division of labor + centralized decision-making. In the analysis layer this study sets three expert agents which represent different investment philosophy: a Value Investing Agent, Growth Investing Agent and Macro Analysis Agent, and each is charged with the responsibility of analysing the market and the different components of the market and stocks in depth. This study puts in place at the decision-making tier a CIO Agent who dynamically incorporates the opinions of experts and implements definitive risk control and trading decisions. It is a multi-agent, distributed, and multi-layered architecture, which aligns with current advancements in the field of multi-agent financial applications (Yu et al., 2024; Xiao et al., 2024), but emphasizes the construction of a more realistic simulation of how a real-world investment research team makes decision-based choices instead of depending on reinforcement learning to optimize its strategies (Lee et al., 2020).

This study did a long-term backtest (24 years 2000-2024) to test the system and its effectiveness. The main results show that in relation to the single-agent system of benchmark the proposed multi-agent system has exceeding superiority in all essential performance indicators. Its total returns (1482.31% vs. 700.01%) and the annualized returns (11.68% vs. 8.67%) were much more. What is more important is that its risk-adjusted return (Sharpe Ratio: 0.5830 vs. 0.2961) increased almost twice, its maximum drawdown (-48.88% vs. -77.45%) decreased significantly, and the system has managed to turn its Alpha negative into positive (0.0494 vs. -0.0153), which shows its outstanding capacity to produce excess returns regardless of the market.

These empirical findings are supported by three theoretical merits of multi-agent architecture. To begin with, the system is helpful in overcoming the cognitive biases by its intrinsic checks and balances created by various investment views (Grosu, F., 2025). Second, the specialization of highly skilled personnel defines a synergy of analytical depth and breadth that does not simply refer to the weakness of the single model the jack-of-all-trades, master-of-none. Lastly, the dynamic weighting mechanism of CIO agent, provides the system with a potent environmental adaptability enabling it to change its decision-making focus as the market conditions shift wherein the idea is aligned to the reinforcement learning features of multi-agent consideration of risks, which are dynamic.

The major research contributions of this paper are two-fold. Hypothetically, it suggests and confirms a multi-agent LLM investment system which emulates a professional research group. In fact, it affirms the two-fold efficiency of this framework on improving long-term returns and mitigating risk. Further studies may involve more intricate models of agent interaction (debate or negotiation) to enhance the quality of decisions more (Lussange et al., 2021) and automate the process of finding and optimization of investment strategies (Kou et al., 2024). These developments would drive the field of financial modelling to an even more sophisticated paradigm, registering this long-standing request by the upper levels of the finance education community of using advanced models (Jiang et al., 2023).

2. Research Design

2.1. Data Sources

To guarantee analytical impartiality and reproducibility, this study constructed a complete dataset spanning from January 2000 to December 2024 representing one market cycle, combining financial and macroeconomic indicators from various public sources. Among them, asset prices and technical indicators were anonymized historical data of an XETF proxying real world QQQ index fund, including closing prices, trading volumes, and relevant technical indicators at each day. The involved macroeconomic variables, including interest rates, inflation rates, employment information and GDP, were sourced from the Federal Reserve Economic Data system, FRED. In addition, this study also involved the proxy variables of market volatility and investor sentiment, including VIX index and

Consumer Confidence Index. All data were downloaded and analyzed by professional financial analysis software.

In the pre-processing stage, to avoid the model uses certain dates for temporal prediction, this study converted original time stamps into relative monthly serial numbers. To avoid the model dependence or memory bias caused by certain specific assets, this study replaced the original name of asset XETF.

In addition, this study also did multi-level cross-verification on the integrity, consistency and accuracy of data to guarantee the reliability of experimental results.

2.2. Model Design

A design of the operational mechanism of Single-Agent System (Control Group) is as follows. Once the monthly decision is made, the system packages together all the pre-processed market data such as XETF historical prices, key market technical indicators, macroeconomic data briefs, and market sentiment alerts into a single prompt. This alert is then directly put into the Qwen model who is asked to take the profession of a professional quantitative fund manager. And lastly, the model produces an unambiguous trading decision (e.g., BUY: X%, SELL: X% or HOLD) and justification, which is practically mapped in the process of backtesting.

Multi- Agent System Architecture is planned as a bi-leveled structure with the specifics of the division of labor in any type of specialization and centralized decision-making.

Layer 1: Specialist Agents. In this layer, three professional agents will work in parallel: The Value Agent processes valuations and fundamentals according to the Graham-Buffett philosophy. The Growth Agent measures momentum and growth ratios using the Fisher-Cathie Wood model. The Macro Agent tracks the indicators of the systemic risk like the yield curve and the VIX (volatility index). Both agents give a standardized report after which there are multi-dimensional scores (1-10) and suggestions.

Layer 2: CIO Agent. The specialist agents have their outputs combined with this agent. It integratively rates the opinion of individual experts depending on the current market situation, it produces a particular amount of a position based on scoring syntax, and it imposes hard-tuned risk constraints, such as a trade size constraint of 10 percent, a total part size constraint of 90 percent, and a naked shorting ban which guarantees strong and executable decisions.

3. Experimental Results and Performance Comparison

Based on a backtest conducted over the full market cycle from January 2000 to December 2024, the two strategies exhibited significant differences in performance.

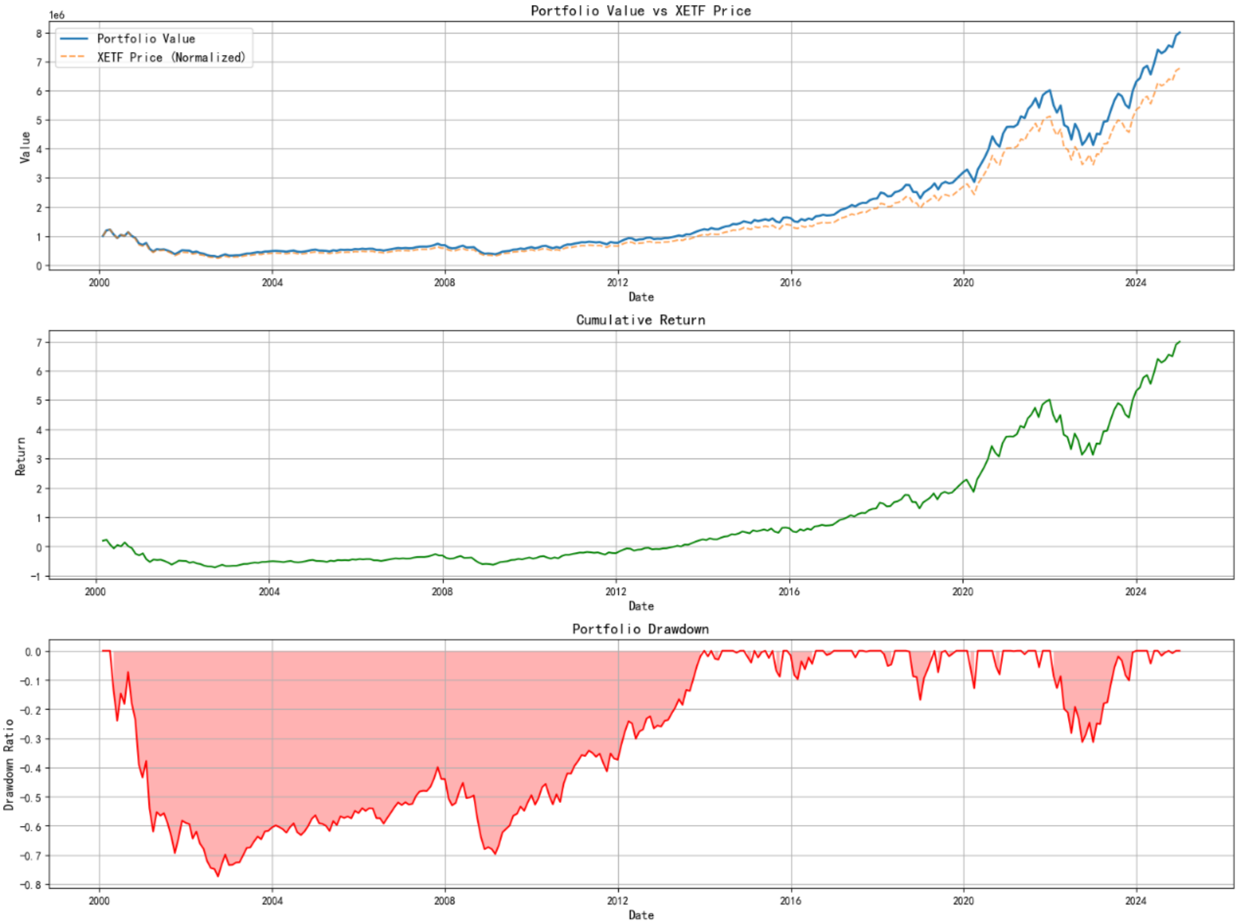

Figure 1: Single-Agent backtest results

The performance of single-agent system is well shown in Figure 1. Its net worth portfolio (blue solid line) followed the model on the market (XETF) (orange dashed line) that is the Nasdaq QQQ most of the 24-year period. The two lines were virtually moving in tandem especially prior to 2016. This means that the returns made by this system were mainly due to passive exposure to beta to the larger market without being able to generate any excess returns (alpha). This finding is supported by the values of Table 1: its Beta value is quite high (0.97), which moved almost in the same direction as the market, but its Alpha (-0.0153) showed that the strategy did not over perform the benchmark. Though, with the tech stock cheap market post-2020, the portfolio value increased, the cumulative return curve slope (green line) is much less steep than Figure 2, and the overall cumulative returns (700.01%) was only comparable to the benchmark growth. Moreover, the drawdown chart below indicates that it experienced a devastating maximum draw down of -77.45% in 2000 during the dot-com bubble and the 2008 financial crisis which was a sign of poor risk-management hence leading to low Sharpe ratio (0.2961). Overall, the single agent system was rather a market tracker example, which did not demonstrate active management capabilities to beat the market.

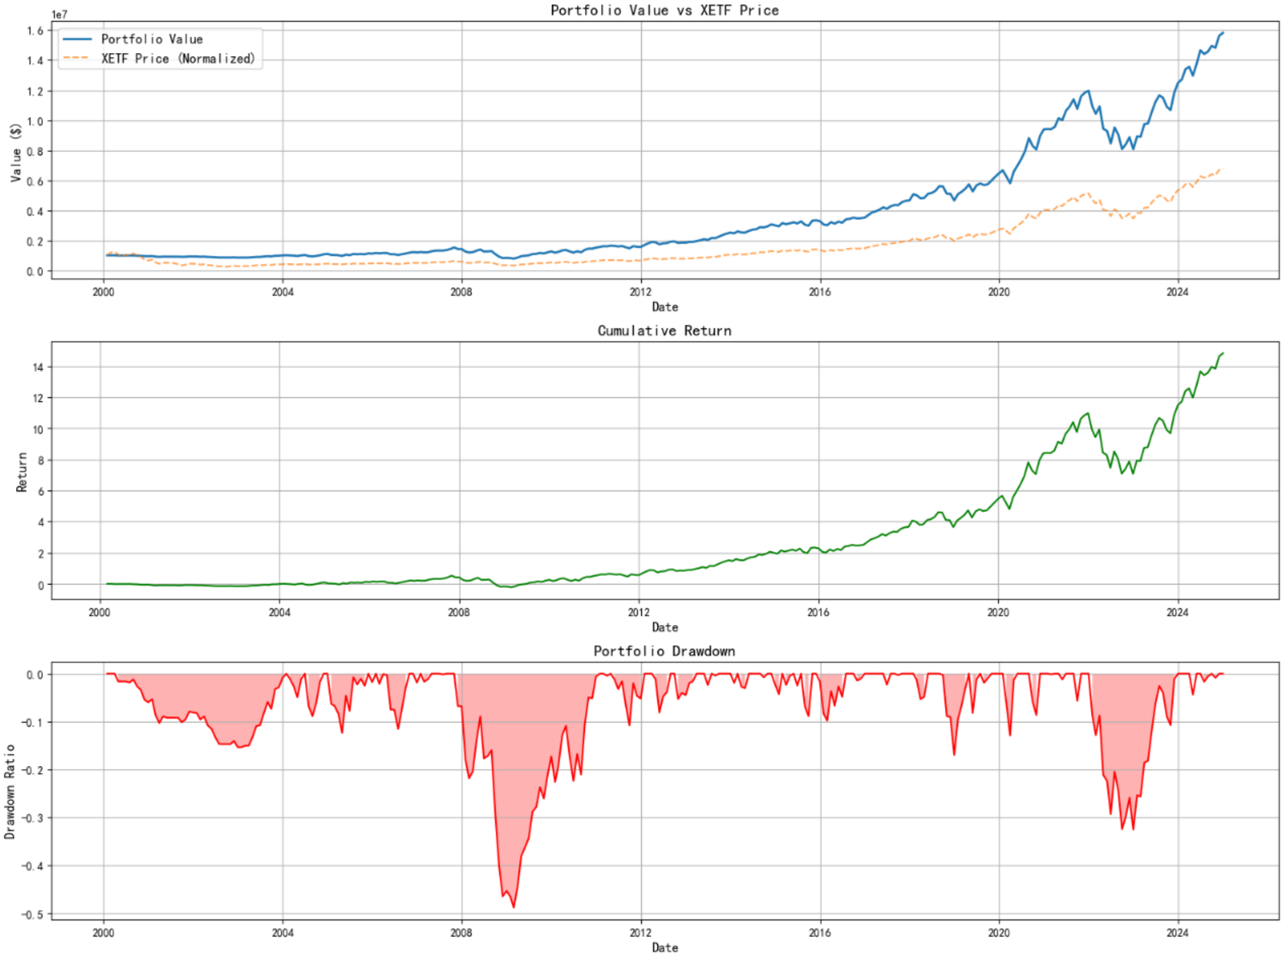

Figure 2: Multi-Agent backtest results

Figure 2 is quite different. It displays the excellent performance of the multi-agent system. Its portfolio net worth (blue line) started to deviate from the benchmark (orange line) around 2010. It experienced explosive growth after 2016 and it is far superior to the benchmark. It is obvious that the multi-agent system possesses a great ability to harvest excess returns. This is also supported by the total return (1482.31%) and the positive Alpha (0.0494) displayed in Table 1. The steep part of the cumulative return curve (green line) in the later years indicates that the system exhibits a powerful compounding effect. What is more important, the system harvested high returns while controlling risks well. Although the system could not escape drawdowns in years of systemic market risks (marked by the red area below), its maximum drawdown was -48.88%, which was much smaller than that of the single-agent system. This result indicates that the internal checks and balances among the three perspectives and strict risk management by the CIO greatly mitigated downward shocks. Finally, the system obtained a Sharpe ratio that was close to a doubling (0.5830), and a substantially lower Beta (0.5614). The system won the battle in both more independent strategy and higher investment efficiency. The tremendous victory was its total victory in risk-adjusted returns.

Table 1: Performance metrics comparison of Single-Agent vs. Multi-Agent systems

Performance Metric	Single-Agent System	Multi-Agent System	Improvement
Total Return	700.01%	1482.31%	+111.8%
Annualized Return	8.67%	11.68%	+34.7%
Sharpe Ratio	0.2961	0.5830	+96.9%
Sortino Ratio	0.4041	0.8328	+106.1%
Maximum Drawdown	-77.45%	-48.88%	+36.9%
Alpha	-0.0153	0.0494	Negative to Positive
Beta	0.9727	0.5614	-42.3%

As it is evident in Table 1, the multi-agent system had a significantly better performance over the 24 years backtesting period than the single-agent system. It had a total return of 1482.31 which was better by 111.8 per cent than the 700.01 of the sole agent. This growth has been remarkable 34.7 percent a year, whereas the annualized return grew 11.68 percent compared with the prior growth of 8.67 percent, which is entirely justified by the major benefit of the multi-agent architecture in maximizing the profit potential. At the same time, there was a significant enhancement of the risk control capabilities in the system. Maximum drawdown reduced to -48.88 as compared to -77.45% in the single system, which represents a 36.9% decrease in the exposure to risk, which indicates the physical influence of the multi-layered risk control system in eliminating extreme losses.

It was also through a tremendous optimization of the efficiency of the risk-adjusted returns. Sharpe ratio has increased by 96.9 percent, 0.2961 to 0.5830 whereas the Sortino ratio increased by 106.1 percent; 0.4041 to 0.8328. This shows that the multi-agent system had almost increased its effectiveness by a factor of two in the returns that it earns per unit of total or downside risk assumed. Importantly, the system made a breakthrough in excess generation returns, and the Alpha component of it is changed to a negative -0.0153 and is greatly positive, 0.0494. This shows that it was capable of always beating the benchmark and getting positive returns regardless of the trend in the market. Moreover, the system became much less responsive to the general market variations, as the Beta value of the system declined to 0.5614 (a decrease of 42.3 percent). This implies a significantly reduced exposure to systemic risk of the portfolio which translates into more independence and stability of the strategy.

4. Discussion

4.1. Core Advantages of the Multi-Agent System

The overall performance of the multi-agent system demonstrated in this paper is much more effective, which relies on the fact that the framework of the intrinsic system is sensible. One, the cognitive bias system is developed subconsciously because the functions of multiple participants in a process are separated. Value oriented agent is conservative in the valuation and margin of safety, growth-oriented agent is momentum oriented and growth potential oriented and the macro view agent is ever keen to see the systemic risks and market sentiment shifts. The simultaneous establishment and resultant integrating of these variegated views has a winning effect on the ability to neutralize the decision-making error of the overconfidence of one model, anchoring effects, or impacts of illusion of information. It is in this light that the multi-agent system was able to make its Alpha positive and significantly reduce its maximum draw- down.

Second, the organized division of labor is a specialized one and enhances the level and scope of the whole analysis. The agent analysts are concerned with such critical variables in its respective investment philosophy that is not subjected to a facade of thought that could have been brought about by an agent analyst attempting to capture all the dimensions in a restricted context. The CIO agent will then proceed with the global integration whereby special knowledge of various domains will be preserved and the final decision will be made to have a coherent set of decisions. It is a synergistic process that challenges the inherent vulnerability of the effectiveness of the given agent that could be either: the mile wide and an inch deep, or: the inch wide and a mile deep.

Finally, environmental flexibility is among the significant supporting pillars of the power of the system. The dynamic mechanism of the weight allocation that is dynamic enables CIO to transform dynamically the value of different perspectives under varying conditions of the end result of the market condition. It, as an example, automatically puts more weight on the macro agent in 2020-2022 on the high volatility of the monetary policy and sharp movements, causing timely risk detection and defensive de-risking flows and, thus, the downside losses under control.

4.2. Effectiveness of the Risk Control Mechanism

The results of the risk control also indicate the effectiveness of its multi-layered mechanisms, acting in a synergetic way. The value agent with its demanding validation of the level of valuation established a restraint on the urge of pursuing rallies at the high-market levels. Risk thresholds, like those determined by the VIX, led the macro agent to provide warnings and initiate changes in positions in the 2022 cycle of Federal Reserve interest rate increases. Moreover, the CIO layer imposed hard-coded constraints, such as the size of individual trades, overall exposure to equity as well as the ban on naked short-selling, which in essence inhibited non-rational or non-compliant operations. These three defense lines operated together and allowed the system to have discipline and remain in order even in the extreme market conditions which eventually led to a 36.9 percentage gain on maximum drawdown and such large Beta decrease.

4.3. Practical Value and Application Scenarios

The multi-agent system suggested in the given study is useful in real-life finance applications to a large extent. The system can also meet the dual goal of high returns and low drawdowns and has a Sharpe ratio that is almost twice that of the single agent system and it is in this respect that is very much in line with the fundamental mandate of institutional investors of getting higher risk-adjusted returns. The system also has the effect of correcting human emotional interference in the investment decision making process by automating and standardizing it, thus ensuring a disciplined implementation of the process and especially when the market becomes very volatile or extreme thus greatly alleviating operational risk compared to behavioral biasness. In addition, the framework has a strong scalable research ability, it is capable of conducting the complete level of assessment in the market with one man size research team per month of the hand-calculated research which establishes significant prospects to be replicated in a business setting.

Regarding the application scenarios, the system can be used as a robo-advisor engine that dynamically changes the weights of various agents depending on the riskiness of a client to provide him or her with asset allocation recommendations. In the case of asset management companies, it may act as an artificial intelligence analyst group to fund managers and quickly generate multi-dimensional investment theses that are structured to aid human decision making. The framework may be also incorporated as an enhancement module into the current quantitative systems, and the semantic knowledge and logical reasoning of the LLMs may be applied to the traditional factor models to address the weakness of the exclusive statistical systems in elucidating underlying logic. Also, macro and market sentiment agent systemic risk warnings can be used to assist tactical asset allocation and dynamic risk control modifications on a portfolio level and increase the overall portfolio flexibility and adaptability.

5. Conclusion

It performed a systematic backtest and comparison of single agent and multi agent large language models in stock investment decision making on 24 years of past data (January 2000 to December 2024). The findings provide a conclusive argument that the multi-agent system is by far better compared to the single-agent approach in generation of returns, risk management, and generation of alpha. The system had a total 1482.31% increase which was 111.8 higher than the single agent 700.01. It had an annualised return of 11.68 and Sharpe ratio of 0.5830 which provided almost twofold risk-adjusted performance compared to the single-agent. At the same time, its peak drawdown was actually reduced to -48.88 and its Beta figure narrowed by 42.3 showing increased resistance to risk and reduced exposure to systemic risk of the market. Most importantly, the multi-agent architecture was able to transform its Alpha negative to positive achieving 0.0494 which justifies the fact that it has the capacity to generate excess returns repeatedly irrespective of market trends.

There are three contributions in this paper. First, structurally, it suggests multi-agent investment format relying on a division of labour being specialised, which simulates the interactivity of the collaborative environment of professional research with respect to value, development, and macro perspective of analysis. Second, empirically it confirms the two-fold benefits of this framework in improving returns and risk control with a long-term backtest with full bull and bear market cycles. Lastly, it, in a mechanistic perspective, introduces three fundamental routes of operation, not only cognitive bias cancellation, but also specialized depth synergy, and dynamic environmental adaptation, which serve as theoretical accounts on the success of collective intelligence in a financial decision making.

Although the multi-agent architecture has achieved great success, this research has a number of limitations. A system necessitates simultaneous calls made to several LLM instances, which means that it accounts for significantly more computational resources usage than the single-agent solution; and this might be reduced in the future by model distillation or lightweight inference, or asynchronous scheduling policies. The existing agent interaction model is a comparatively fixed model, in which the CIO performs weighted integration, and is not based on more complicated agent interactions. Future studies may consider dynamic negotiation processes through debate, consensus voting, or reinforcement learning as a way to improve the quality of decisions even more. Also, the monthly rebalance rate is apt to medium- and long-term investment choices; it would need improvement of the system real-time reactiveness and data processing efficiency in case one wants to rebalance it on high frequencies of trading. Future studies will involve the introduction of the reinforcement learning to maximize the weight allocation algorithm, the investigation of the debate and consensus-building game amongst the agents, as well as the extension of the framework to the case of multi-asset classes and global market allocation, and the analysis of its applicability to the situation of higher frequency trading.

In conclusion, the multi-agent architecture provides a new avenue on the faithful utilization of giant language models in sophisticated and stakes-based financial decision-making, which is certain to spearhead the innovations in artificial cleverness in asset management towards more intelligence and solidity.

References

[1] Cao, B., Wang, S., Lin, X., Wu, X., Zhang, H., Ni, L. M., & Guo, J. (2025). From deep learning to LLMs: A survey of AI in quantitative investment. arXiv preprintarXiv:2503.21422.

[2] Ding, H., Li, Y., Wang, J., & Chen, H. (2024). Large language model agent in financial trading: A survey. arXiv preprintarXiv:2408.06361.

[3] Grosu, F. (2025). How can multi-agents AI systems help reduce biases in trading algorithms? Review of International Comparative Management.

[4] Jiang, J., Kelly, B., & Xiu, D. (2023). (Re-)Imag(in)ing price trends. The Journal of Finance, 78(5), 2617–2663.

[5] Kelly, B. (2023). The virtue of complexity in return prediction. The Journal of Finance, 78(4), 1839–1886.

[6] Kou, Z., Yu, H., Luo, J., Peng, J., Li, X., Liu, C., Dai, J., Chen, L., Han, S., & Guo, Y. (2024). Automate strategy finding with LLM in quant investment. arXiv preprintarXiv:2409.06289.

[7] Lee, J., Kim, R., Yi, S.-W., & Kang, J. (2020). MAPS: Multi-agent reinforcement learning-based portfolio management system. In Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI)(pp. 4596–4602).

[8] Li, W. W., Kim, H., Cucuringu, M., & Ma, T. (2025). Can LLM-based financial investing strategies outperform the market in long run? arXiv preprintarXiv:2505.07078.

[9] Lussange, J., et al. (2021). Multi-agent reinforcement learning for market microstructure statistical inference. Journal of Financial Data Science, 3(4), 112–127.

[10] Malberg, S., Poletukhin, R., Schuster, C. M., & Groh, G. (2024). A comprehensive evaluation of cognitive biases in LLMs. arXiv preprint.

[11] Wall, L. F., & Leitner, S. (2021). Agent-based computational economics in management accounting research: Opportunities and difficulties. Journal of Management Accounting Research, 33(3), 189–212.

[12] Xiao, Y., Sun, E., Luo, D., & Wang, W. (2024). TradingAgents: Multi-agents LLM financial trading framework. arXiv preprint.

[13] Yu, Y., Yao, Z., Li, H., Deng, Z., Cao, Y., Chen, Z., et al. (2024). FINCON: A synthesized LLM multi-agent system with conceptual verbal reinforcement for enhanced financial decision making. In Advances in Neural Information Processing Systems (NeurIPS).