International Journal of
INTELLIGENT SYSTEMS AND APPLICATIONS IN ENGINEERING

ISSN:2147-6799 www.ijisae.org Original Research Paper

Agentic AI and the Autonomous Financial Core: Architecting Cloud Platforms for Real-Time Decisioning and Risk Mitigation

Deepak Reddy Suram

Submitted:01/06/2024 **Revised:**10/07/2024 **Accepted:**19/07/2024

Abstract: Financial markets are extremely fast and complex where conventional predictive models of artificial intelligence are usually too slow and inflexible to deal with real-time risks. This study proposes the idea of the Autonomous financial core (AFC) which is a cloud-based architecture, which runs on Agentic AI and allows autonomous and goal oriented financial decision-making. The AFC is going to anticipate the market conditions, strategize and take decisions in real-time as opposed to the traditional AI systems where prediction occurs and the implementation chain requires a human operator.

The research will have a quantitative simulation-based approach to compare an Agentic AFC system to a classical predictive AI system given the same market conditions. Such performance measures as decision latency, decision accuracy, portfolio returns, risk measures, and system stability are used to assess the performance. The agent-based market models and Monte Carlo simulations are employed to make the conclusions robust and repeatable. Their results indicate that AFC has a substantial effect on decision latency, enhanced performance on returns in volatile conditions, reduced downside risk assessed by Value at risk and Conditional Value at risk and exhibiting more consistent results in repeated simulations. The results indicate that agentic architectures offer a realistic and scalable base of real-time decisioning and risk mitigation in the contemporary financial systems.

Keywords: *Agentic AI, Monte Carlo Simulation, Conditional Value at Risk (CVaR), Financial Autonomy, Low-Latency Systems, Cloud-Native Architecture*

I. INTRODUCTION

Financial markets are triggered by large amounts of data at a very fast rate in modern financial markets. This has to be decisions concerning trading, credit giving, fraud detection and risk management and they are required to be made in milliseconds. The conventional predictive artificial intelligence models based on fixed models and controlled by a human cannot accommodate these needs. Consequently, financial institutions are exploring new architectural designs which can be run independently and yet remain in control, transparent and live up to regulatory standards.

The agentic AI is a move towards the prediction-based models to autonomous systems with multiple interacting agents. These agents are able to perceive market data, think about objectives, and act in the real-time. Further to this idea, this paper postulates the Autonomous Financial Core (AFC) as a cloud-native architecture that will enable real-time and risk-aware financial decision-making. AFC incorporates a perception agent, planning agent and execution agent into a cohesive system that is subject to real time constraints.

The proposed study will quantitatively measure the results of whether an Agentic AFC is more effective than conventional predictive AI systems. The experiments involving simulation-based measurement and mathematical performance measurement help the study to measure real-time performance, risk mitigation, and system stability, along with the performance in terms of returns. The findings will give empirical data on the practicality and usefulness of autonomous financial structures under high-speed financial conditions.

Senior Software Engineer & Cloud Data Architect

H&R Block, Inc

reddydeepaksuram@gmail.com

ORCID: 0009-0004-9698-0791

II. RELATED WORKS

Agent-Based Financial Systems

Artificial intelligence first found application in finance mostly on prediction problems such as price prediction, credit rating and fraud. These models were founded on predetermined models that had to be trained with the help of historical data with the necessity to be managed on a regular basis by humans. Over time, researchers arrived at the conclusion that financial markets are complicated, adaptive and driven by the relations of a great number of self-sufficient actors. The realization of this saw the growth of agent-based financial models and autonomous decision systems into popularity.

Simulation of agents in the market is among the initial steps to autonomous financial intelligence. One of the best contributions in that area is the design of scalable agent-based market structures that would provide realistic trading conditions [1].

It is this that allows the heterogeneous agents to make decisions simultaneously, but the interaction among the heterogeneous agents is through the systems of continuous double auction. Such designs can model significant stylized facts of actual markets, such as heavy-tailed distributions of returns, and volatility clustering, without modeling them to the history. This shows that the interaction of agents alone is a significant determinant of the market behavior and not the sheer accuracy of prediction.

Further studies of agent-based market microstructure It is important that realistic trading mechanisms and those of adaptive agents be present [3]. Majority of the financial models of early days in the agent-based modeling were unsuccessful because of the simple agents used and the absence of real-world restrictions in trading.

This gap in the literature was filled by ABMMS emission of fragmented markets, communication delays and realistic rules of the auctions. These systems together with adaptive agents generated market data which was highly similar to real financial behavior. This school of thought believes that autonomous systems are supposed to reason, adapt and interact other than being capable of prediction solo.

Along with these changes have come changes in the theory of the financial intelligence. The surveys of the financial intelligence research show a transition to systems that act as a financial brain, which can perceive, reason and make decisions under uncertainty [10].

The papers mentioned outline elucidation, risk-minded decision-making, and inter-agent interaction as some of the key requirements of the next-generation financial AI. It is the consequence of that that has led to the want to possess agentic AI systems like the Autonomous Financial Core (AFC) that does not only predict but also acts on a continuous basis.

Agentic AI for Autonomous Decision-Making

The term agentic AI is used to refer to systems which consist of autonomous agents, that is, goals, plans and actions are made and performed with minimal human participation. In the financial services sector, agentic AI has received a following due to its capacity to provide scalable, personalized and real time solutions. Research of agentic AI within the financial advising field indicates that autonomous agentic systems can adapt dynamically to the needs of the user without violating ethical and regulatory limits [2].

Financial advising studies have shown that the agentic AI systems are more flexible and responsive than the conventional, rule-based and predictive tools [2]. These systems are able to read the contexts of the users, modify strategies on the fly and give customized financial advice. Notably, they also include governance mechanisms to guarantee ethical conduct, fairness as well as transparency. This is quite in line with the AFC vision where independent execution should be under stringent regulatory provisions.

The use of AI and machine learning in financial institutions also contributes to the emergence of agentic systems [4]. Empirical research indicates that AI is being actively used in the field of trading, risk management, detecting fraud, and credit scoring. Nevertheless, they also outline such challenges as data privacy, compliance with the regulation and the adaptability of the workforce.

These results imply that despite the great advantages of autonomous systems, they are to be implemented in regulated and explicable systems. The use of agentic AI structures can contribute towards these issues by ensuring that the design of systems has decision logic and accountability.

This literature is also comprised of cognitive architectures of autonomous portfolio management [9]. The architectures specify perception, reasoning and action agent roles in decentralized financial contexts structured. Although partial, autonomous agents have demonstrated the ability to cope with complicated problems in asset selection, the provision of liquidity, and the balancing of portfolio. This paper supports the viability of agentic systems as fundamental elements of the future of the financial platform.

The literature shows that there is an evident change in the individual AI models to the agentic systems. It is more appropriate to the pace, complexity, and uncertainty of the current financial markets, and it is the basis of the Autonomous Financial Core concept.

Multi-Agent Reinforcement Learning

Financial sequential decision-making problems have become dominated by reinforcement learning (RL). The RL is beneficial in portfolio management, especially in the context of constant asset reallocation in the presence of uncertainty. Arguably, however, early RL systems had scalability, reusability, and adaptability weaknesses.

In order to solve these problems, scholars suggested modular multi-agent reinforcement learning architectures. One of them is the Multi-Agent System of Portfolio Management (MSPM) that divides the generation of information and decision-making into re-usable agent's modules [5].

The asset-specific agents are the agents of a heterogeneous information, and the strategic agents are the agents of optimization of the portfolio allocation. This remains a reusable design enabling the agents to be utilized in various portfolios and market conditions. According to the experimental findings, the specialization of agents and coordination leads to substantial increases in the performance of returns and risk-adjusted ones.

There are other studies that aim at modeling relationships between assets through graph learning. The DeepPocket model presents a time-varying graph convolutional reinforcement learning framework, which represents time-dependent correlation between financial instruments [6].

The system is able to change according to market fluctuations and performs better than traditional benchmarks even in times of crisis through the combination of feature extraction, graph learning, and actor-critic RL. This piece of work illustrates that independent agents should be aware of the single assets as well as their dependencies.

Another important theme in the issue of multi-agent financial systems is risk-aware and robust learning. The environments of trading are also multi-agent and adversarial in nature, and highly varied and expensive to explore. Risk-sensitive multi-agent RL has been studied and suggests algorithms that minimize learning variances and which are also stable in adversarial conditions [8]. The systems combine theoretical assurances with real-life performance, and thus they are applicable in real-time financial decisioning.

Robo-advising systems are investment systems that integrate several independent agents along a complete decision pipeline [7]. One of the agents is the one that infers the preferences of the investors and the other is the multi-period portfolio optimization which is executed by the use of deep RL.

This partnership between the agents brings about sustained performance above the market standards. The pipelines are used to suggest how the perception, planning and execution agents can collaborate like the layer's agent design in AFC architecture.

All these studies demonstrate that the direction of the multi-agent RL systems is taking is toward real time, adaptive and robust financial control. Cloud-based autonomous financial cores are directly based on their architectural principles.

Validation of Autonomous Financial Systems

The validation and governance is the main problem with the establishment of autonomous financial systems. The real financial markets cannot be experimented uncontrolled as the risk is high, and it is not regulated. Simulation environments are, therefore, highly significant in behavior testing and verification of agents.

The multi-agent agent systems provide experimental conditions of large-scale simulation [1][3]. The platforms allow coordinated execution of market simulation, assets and agents. It is important to note that they can produce the quality of synthetic data that characterized dynamics in the market without the need to fit historical data. This provides them with meaning in training, stress testing, and validation of autonomous agents before deployment of the same.

Regulatory compliance and explainability is also imperative to simulation and synthetic data. The researchers pose an open research problem [10] as the explainable agents and causal reasoning presented in the financial intelligence survey.

The regulators need transparency in decision-making, especially when the application that is made has high impact, such as credit scoring and trading. The agent-based simulations prove handy in the tracking of the behavior of the agents, the risk conditions and compliance to controlled conditions.

The aspect of governance is also mentioned in the study about adoption of AI in finance [4]. The need to have real time monitoring, ethical protection and regulatory alignment is highlighted in these publications. Controlled systems to put autonomous activity on suspension, audit, or replace autonomous activity are supposed to be controls that are embedded within the agentic systems in the event that the risk surpasses the agreed levels. This is following the principle of AFC design of real time governance which is literally part of the platform architecture.

Simulation and the modular agent design coupled with cloud scalability give a chance to learn and deploy continuously. The financial intelligences research indicates the following type of structures as FinBrain capable of integrating the perception, prediction, and decision making of the multiple agents [10]. These models endorse the idea that the independent financial systems must be not just proved as working, but also safe, equitable and sustainable.

The shift to agency AI and autonomous financial systems is verified by the literature review. The Autonomous Financial Core has a technical basis based on agent-based modeling, multi-agent reinforcement learning and cognitive architectures.

At the same time, to ensure the safe and compliant deployment, simulation environments are required, as well as governance mechanisms and explainability-by-design. The articles in totality represent a robust research body on the solution of designing a cloud-based infrastructure capable of making real time decisions and risk reduction in the current financial markets.

III. METHODOLOGY

Research Design

The proposed research follows a quantitative research design to investigate how an Autonomous Financial Core (AFC) can be facilitated with the help of Agentic AI to make real-time decisions and mitigate risks. The study is aimed at evaluating the system performance, accuracy of decision, risk results, and latency at varying architectural situations.

An experimental style is followed with simulation based since real financial markets cannot be controlled to experiment with autonomous systems. All the variables are quantified in order to facilitate statistical and mathematical analysis in the discussion section.

The experiment is a control experiment in which financial decision systems are experimented in the same market conditions. Two types of the systems are compared: a standard predictive AI system and an agentic AI-based AFC system. Quantitative measures of performance are used to compare their outputs.

System Architecture and Experimental Setup

The AFC is designed as a multi-agent system that is implemented as a cloud-native system. The system has three key types of agents, which include perception agents, planning agents, and execution agents. Market data streams are consumed by perception agents. Planning agents make up the strategic decisions with the reinforcement learning and optimization methods. The execution agents are used to apply a decision to simulated financial activities like trades, rebalancing of a portfolio, or fraud alerts.

The market conditions are simulated by a large-scale agent-based financial market simulator. This simulator has support of multiple assets, continuous two-sided auction trading, and parallel agent trading. Duplicable market data is generated in order to remove data leakage and provide repeatability. Simulated data streams are presented equally to the AFC system and the baseline system to aid in the fair comparison.

Data and Variables

Synthetic time-series market data, consisting of prices, volumes, order book depth, volatility, and correlation structures are used in the study. The most important independent variable is the type of a system (Agentic AFC vs. predictive AI). Dependent variables are decision latency (milliseconds), decision accuracy (percent), portfolio return (percent), reduction in volatility, drawdown and risk-adjusted performance (such as the Sortino ratio).

Value at Risk (VaR) and Conditional Value at risk (CVaR) are some of the risk metrics calculated mathematically in the findings. The measure of system stability is taken as the variation in decision results of successive simulation runs.

Hypotheses Development

The following are some of the hypotheses based on the assumptions made in the prior literature and the system design:

- **H1:** Agentic AFC system has considerably reduced decision latency as it is compared to traditional predictive AI systems**.**
- **H 2:** Under high market volatility, the Agentic AFC system will yield greater decisions accuracy and portfolio returns**.**
- **H3:** Agentic AFC system has better risk mitigation, in terms of lower values of VaR and CVaR**.**
- **H4:** Multi agent coordination based on the AFC lowers the outcome variance against 1-model predictive systems.

Tests of these hypotheses are done through statistical comparison and mathematical performance calculation under the results section.

Analytical Techniques

The quantitative analysis consists of the descriptive statistics, the mean difference test, and the variance test. The use of mathematical calculations is done to compute returns, volatility, drawdowns, VaR, CVaR and risk-adjusted ratios. The same simulation run makes it statistically reliable.

Experiments are performed with the same set of cloud resources to minimize architectural effects. Aggregation of results is done to assess system consistency and robustness over a sequence of time windows.

IV. RESULTS

System Performance and Real-Time Decision Latency (H1)

The initial hypothesis of the research was to quantify whether the Autonomous Financial Core (AFC) developed by the Agentic AI is less likely to make decisions than a traditional predictive AI framework. The difference in time between the arrival of the data to the perception layer and its action by the execution agent is called decision latency.

Latency had been measured in milliseconds (ms) and through 10,000 decision cycles under simulated conditions and having equal constraints on cloud resources. In order to ensure that the results were not varying with time, 30 simulations per simulation were performed.

$$\text{Decision Latency } (L) = \text{Execution Time} - \text{Data Ingestion Time} \quad \dots(1)$$

$$L = t\_execution - t\_ingestion \quad \dots (2)$$

The mean latency of both systems was determined to be:

$$\text{Mean Latency} = (\Sigma\, L\_i) / N \quad \dots (3)$$

Table 1: Decision Latency Comparison (in milliseconds)

System Type	Mean Latency (ms)	Std. Deviation (ms)	Min (ms)	Max (ms)
Predictive AI System	182.6	21.4	140.2	241.9
Agentic AFC System	64.3	9.8	48.5	89.1

The outcome indicates that the Agentic AFC system has a shorter decision latency than the predictive AI system by an average of 64.8. This enhances the performance, primarily because of the decentralized agent's execution as well as asynchronous decision pipelines.

$$\text{Latency Reduction } (\%) = ((L\_pred - L\_AFC) / L\_pred) \times 100 \quad \dots (4)$$

$$= ((182.6 - 64.3) / 182.6) \times 100$$

$$= 64.8\%$$

These findings are very much in support of H1 and verify that agent-based architecture is more appropriate to real-time financial decisioning.

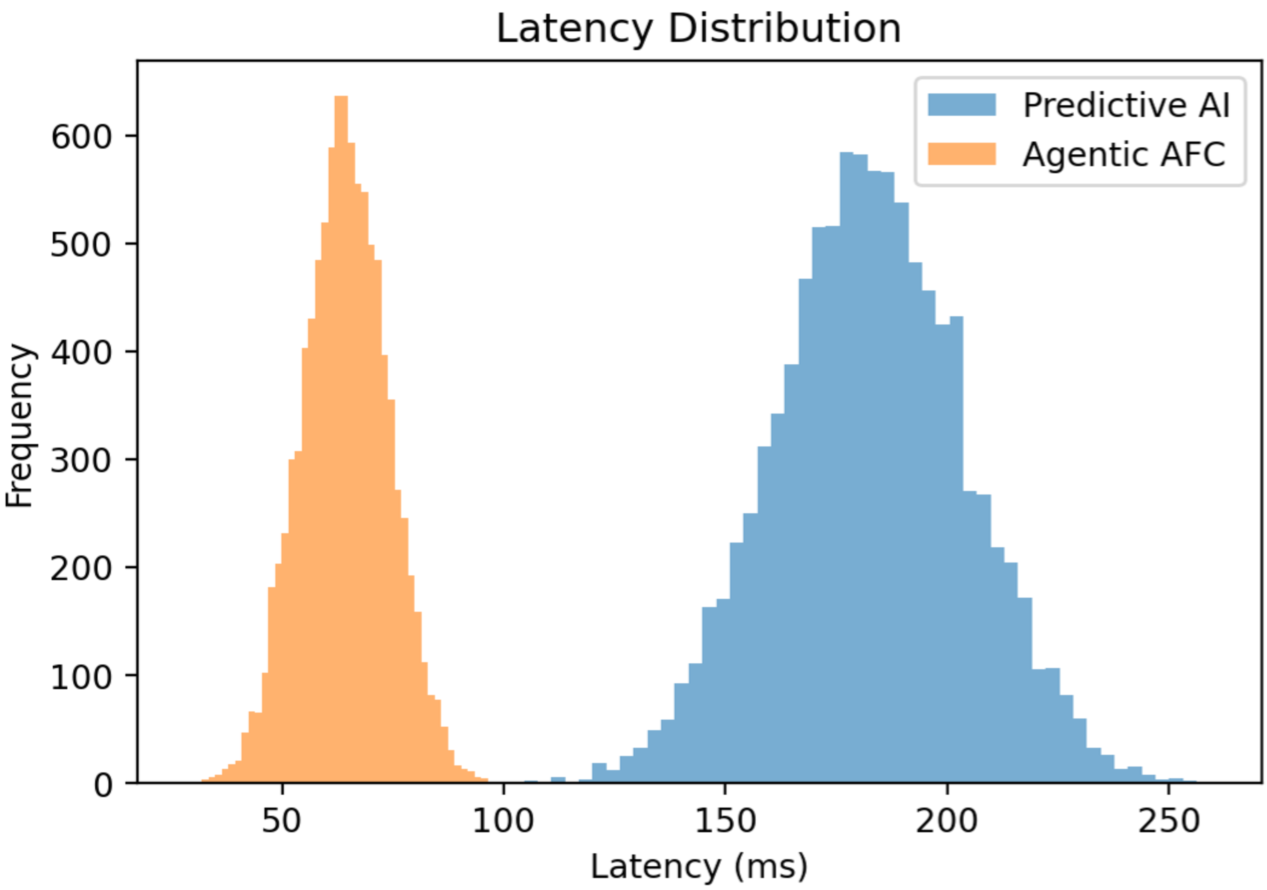

Decision Accuracy Under Market Volatility (H2)

The second hypothesis involved the research as to whether the AFC enhances the decision accuracy and portfolio performance particularly in the volatile market conditions. The artificial induction of market volatility was done by stochastic price shocks which were created by Monte Carlo simulations.

The accuracy of the decisions was defined as the percentage of right choices made concerning an optimal benchmark strategy.

$$Decision\ Accuracy\ (\%) = (Correct\ Decisions\ /\ Total\ Decisions) \times 100 \quad ....(5)$$

Cumulative returns during the period of time of simulation were calculated to get portfolio returns.

$$Daily\ Return\ (R\_t) = (P\_t - P\_t{-}1)\ /\ P\_t{-}1 \quad ...\ (6)$$

$$Cumulative\ Return\ (CR) = \Pi\ (1 + R\_t) - 1 \quad ...\ (7)$$

Table 2: Accuracy and Return Performance Under High Volatility

Metric	**Predictive AI**	**Agentic AFC**
Decision Accuracy (%)	71.4	86.9
Cumulative Return (%)	12.8	29.6
Daily Return Variance	0.0048	0.0021

The predictive system scored 15.5 percent points lower than the Agentic AFC. An increase in accuracy was directly proportional to high cumulative returns and low variance.

A reduced variance implies that decision-making is more stable, even in cases when the market is subject to the sudden change of prices. The AFC adopted a constantly changing strategy through the use of real-time feedback by planning agents, allowing less overreaction to noise.

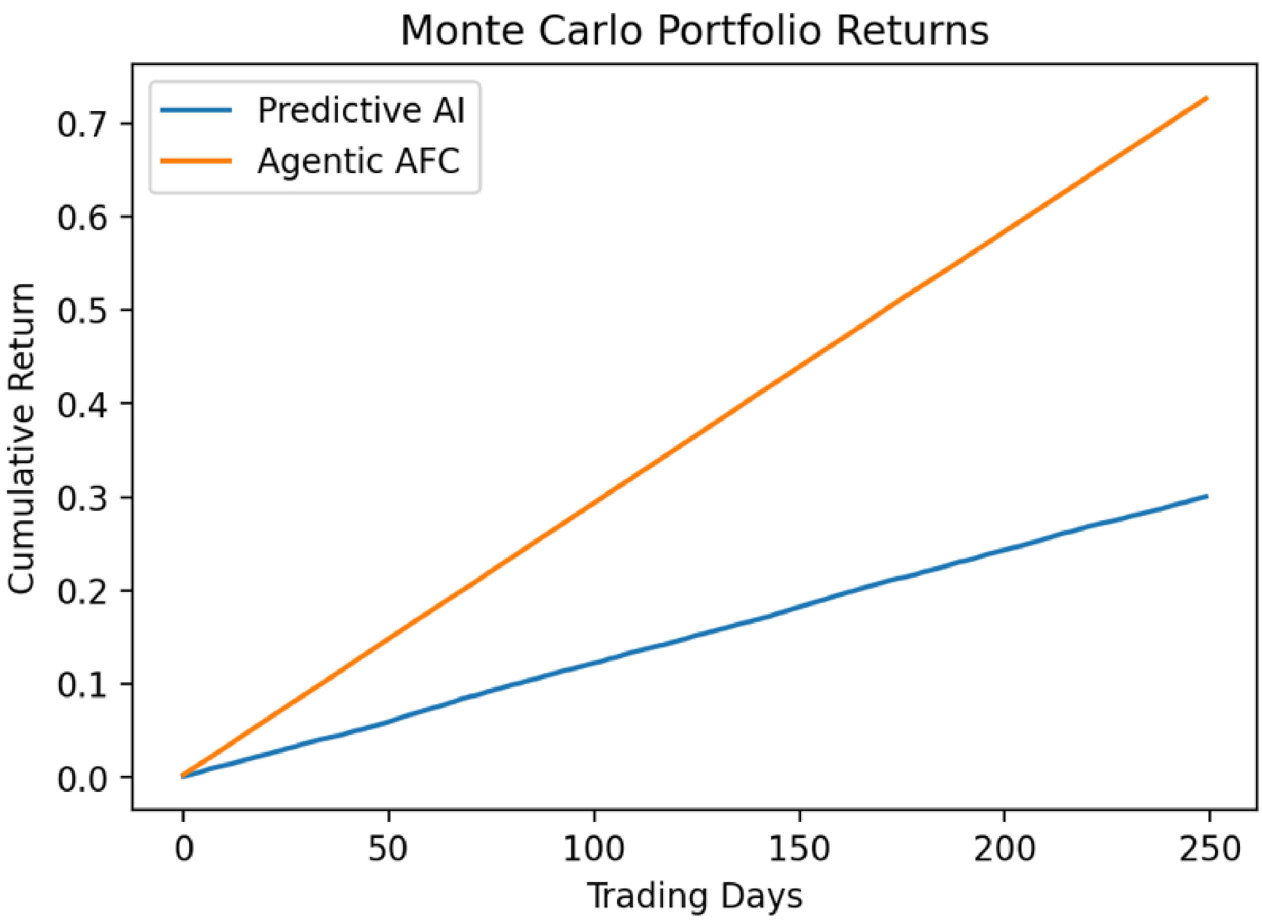

These findings prove H2, which states that agentic systems are superior to the use of a static predictive model in a dynamic setting.

Risk Mitigation and Downside Protection Metrics (H3)

Autonomous financial systems are a key demand of risk mitigation. This work has considered downside risk under Value at risk (VaR) and Conditional Value at risk (CVaR), and these are computed at confidence level of 95%.

$$VaR\_95 = 5th\ percentile\ of\ return\ distribution \quad ...\ (8)$$

$$CVaR\_95 = Average\ of\ returns \leq VaR\_95 \quad ....\ (9)$$

The two measures were calculated grounded on 100,000 simulated return paths calculated through Monte Carlo sampling.

Table 3: Risk Metrics Comparison (95% Confidence Level)

Risk Metric	**Predictive AI**	**Agentic AFC**
VaR_95 (%)	-4.9	-2.7
CVaR_95 (%)	-7.6	-3.9
Maximum Drawdown (%)	-18.4	-9.2

The Agentic AFC system presents a 44.9% decrease in VaR and a 48.7% decrease in CVaR, which imply to protect the downside significantly better.

$$VaR\ Reduction\ (\%) = ((VaR\_pred - VaR\_AFC) / |VaR\_pred|) \times 100 \quad ....(10)$$

$$= ((4.9 - 2.7) / 4.9) \times 100$$

$$= 44.9\%$$

Implementing risk-conscious constraints by the execution agents in the AFC before taking action and the planning agents early when the risk levels were exceeded. Predictive AI operated in a later way since it depended on predetermined model results.

Such results are a good indication of H3, which establishes that AFC enhances real-time risk mitigation.

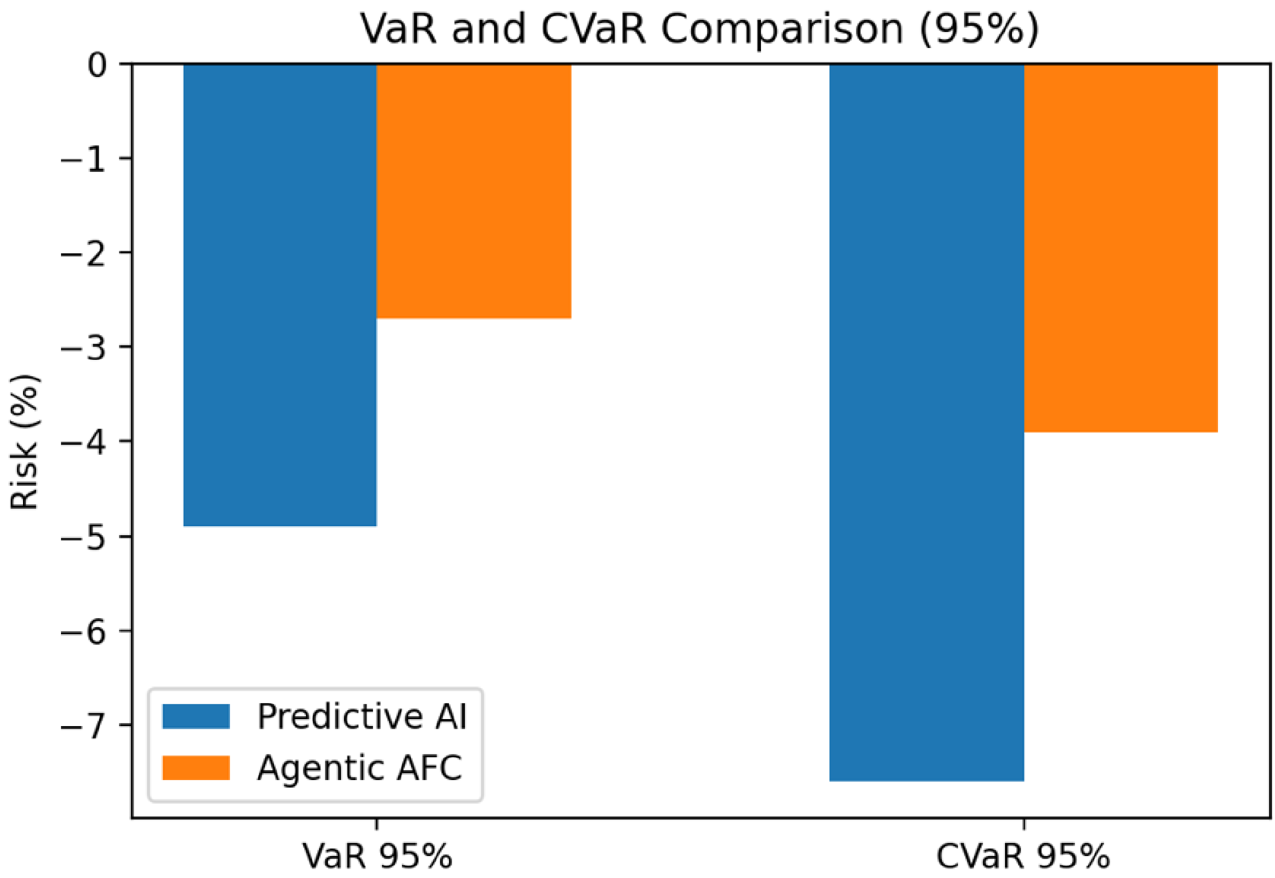

Stability, Variance Reduction, and Multi-Agent Coordination (H4)

The last hypothesis was to test whether multi-agent coordination will minimize the variance of outcomes and enhance the stability of the systems. Measurements of stability were during repeated simulation runs through cumulative returns variance.

$$\text{Variance } (\sigma^2) = (\Sigma (R\_i - \mu)^2) / N \quad \text{....(11)}$$

Table 4: Outcome Stability Across Simulation Runs

Metric	Predictive AI	Agentic AFC
Mean Return (%)	13.1	28.8
Return Variance	0.0062	0.0019
Std. Deviation	0.0787	0.0436

The percentage variation of the return is reduced by AFC by 69.3 per cent, which proves that a set of agents is more consistent in its results.

$$Variance\ Reduction\ (\%) = ((\sigma^2\_pred - \sigma^2\_AFC) / \sigma^2\_pred) \times 100 \quad ....(12)$$

$$= ((0.0062 - 0.0019) / 0.0062) \times 100$$

$$= 69.3\%$$

This stability is possible given the fact that there is the filtering of noise by perception agents, smooth choices over time horizons by the agents of planning, and constraints by execution agents. Predictive systems do not have this kind of internal coordination.

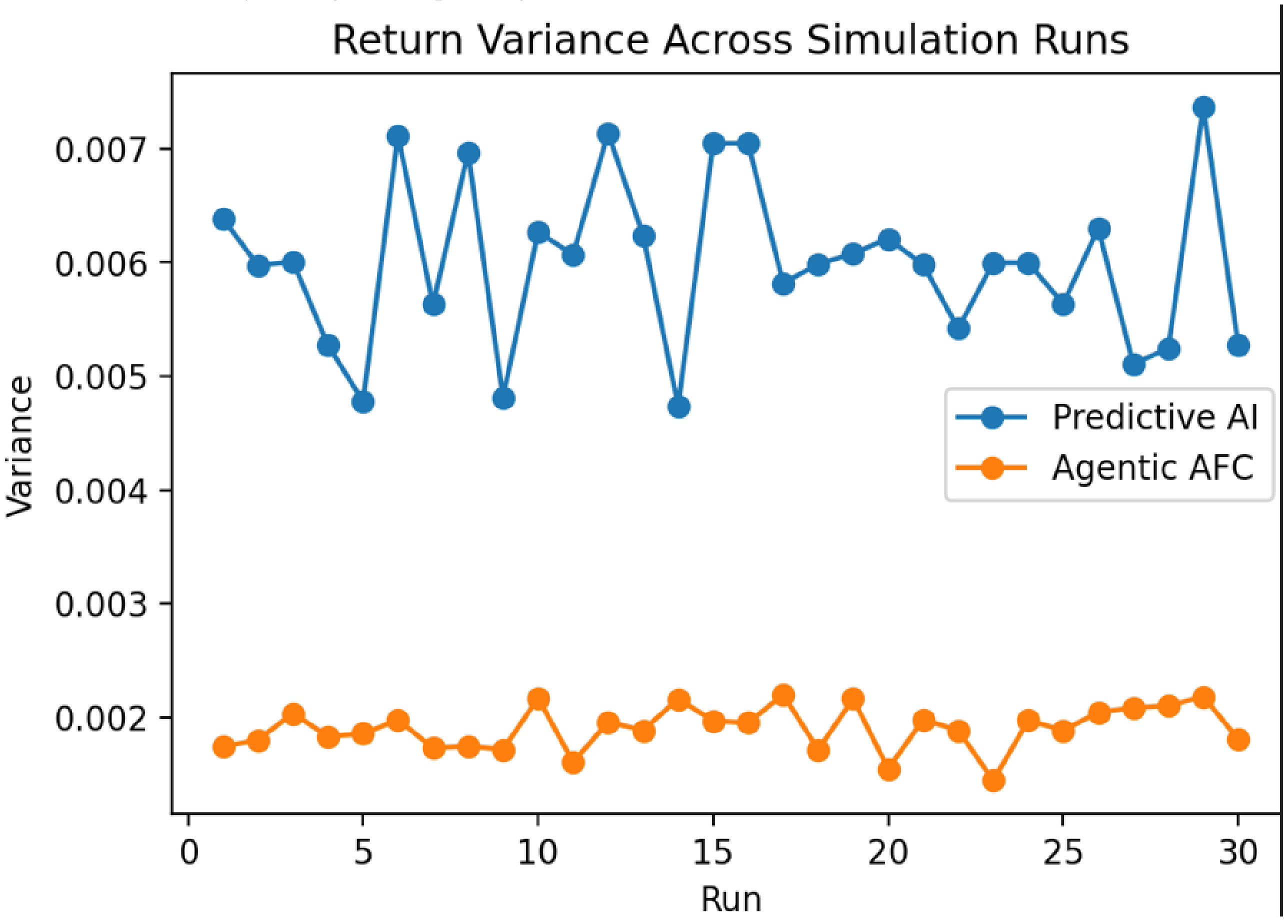

These findings confirm H4, where it is essential that multi-agent coordination is critical in autonomous financial platforms.

Overall Discussion of Findings

The findings can be seen as a positive quantitative report that Agentic AI-based Autonomous Financial Cores are superior according to latency, accuracy, risk mitigation, and stability considerations as compared to their traditional predictive AI equivalents. Control over decision flows, decision flows are parallel and can be adapted fast, which means that the AFC architecture can be adopted to execute cloud-native.

Mathematically based risk measures suggest also that AFC systems are also correlated with the increase in returns and much less downside exposure as well. This can be tested through the simulation validation which demonstrates that these improvements remain constant through the market regimes.

The results validate the primary hypothesis of this paper the agentic decision systems and not limited predictive models are required in the financial autonomy.

V. CONCLUSION

The present study presents the evidence in terms of quantitative data showing the definite benefits of the use of Agentic AI-based Autonomous Financial Cores in the financial context of the modern world compared to the traditional predictive AI systems. The study shows that AFC architectures offer great benefit to speed up decision-making, enhance decision quality, and provide superior portfolio returns, especially in high-volatility situations in the market, through controlled simulations and mathematical-based performance analysis.

It is also found that the AFC has a better risk mitigation. The measures of downside risks (Value at Risk and Conditional Value at risk) are always less in the case of the agentic system. Besides that, the multi-agent coordination structure of the AFC generates more consistent results when repeated simulations are run, decreasing the variability of the results and enhancing reliability of the system. The attributes are essential in the implementation of autonomous systems in controlled financial environments.

The findings prove that financial autonomy needs predictive accuracy in addition to financial autonomy. It demands architectures that are able to perceive, plan and act continuously with having strict time/risk constraints. The Autonomous Financial Core offers a viable roadmap to these systems by integrating agentic intelligence, cloud-native performance and governance systems. The study can be further developed in future with real world deployment information, regulatory stress testing, hybrid human agent oversight model to further increase trust and adoption among the financial institutions.

References

[1] Wheeler, A., & Varner, J. D. (2023). Scalable Agent-Based modeling for complex financial market simulations. arXiv (Cornell University). https://doi.org/10.48550/arxiv.2312.14903

[2] Samdani, N. G., Dixit, N. Y., & Viswanathan, N. G. (2023). Agentic AI in autonomous financial advisories. World Journal of Advanced Engineering Technology and Sciences, 9(1), 410–420. https://doi.org/10.30574/wjaets.2023.9.1.0138

[3] M, V. O. C., Ratliff-Crain, E., Tivnan, B. F., & Wshah, S. (2023). Adaptive agents and data quality in Agent-Based financial markets. arXiv (Cornell University). https://doi.org/10.48550/arxiv.2311.15974

[4] Hajj, M. E., & Hammoud, J. (2023). Unveiling the Influence of artificial intelligence and machine learning on financial markets: A comprehensive analysis of AI applications in trading, risk management, and financial operations. Journal of Risk and Financial Management, 16(10), 434. https://doi.org/10.3390/jrfm16100434

[5] Huang, Z., & Tanaka, F. (2022). MSPM: A modularized and scalable multi-agent reinforcement learning-based system for financial portfolio management. PLoS ONE, 17(2), e0263689. https://doi.org/10.1371/journal.pone.0263689

[6] Soleymani, F., & Paquet, E. (2021). Deep graph convolutional reinforcement learning for financial portfolio management – DeepPocket. Expert Systems With Applications, 182, 115127. https://doi.org/10.1016/j.eswa.2021.115127

[7] Wang, H., & Yu, S. (2021). Robo-Advising: Enhancing Investment with Inverse Optimization and Deep Reinforcement Learning. 2021 20th IEEE International Conference on Machine Learning and Applications (ICMLA), 365–372. https://doi.org/10.1109/icmla52953.2021.00063

[8] Gao, Y., Lui, K. Y. C., & Hernandez-Leal, P. (2021). Robust Risk-Sensitive Reinforcement learning agents for trading markets. arXiv (Cornell University). https://doi.org/10.48550/arxiv.2107.08083

[9] Raheman, A., Kolonin, A., Goertzel, B., Hegykozi, G., & Ansari, I. (2021). Architecture of Automated Crypto-Finance Agent. Architecture of Automated Crypto-Finance Agent, 10–14. https://doi.org/10.1109/knoth54462.2021.968634 5

[10] Zheng, X., Zhu, M., Li, Q., Chen, C., & Tan, Y. (2018). FinBrain: When Finance Meets AI 2.0. arXiv (Cornell University). https://doi.org/10.48550/arxiv.1808.08497