https://doi.org/10.62836/iaet.v3i1.362

Global Science Publishing

IAET

Innovations in Applied Engineering and Technology

https://ojs.sgsci.org/journals/iaet

Article

Generalizable Multi-Agent Framework for Quantitative Trading of US Education Funds

Kaixian Xu [1], Zhaoyan Zhang [2], Alan Wilson [3,*], Yu Qiao [4] Lisang Zhou [5] and Ying Jiang [6]

[1] *Risk & Quant Analytics, BlackRock, Jersey City, NJ 07302, USA*

[2] *Zhongke Zhidao (Beijing) Technology Co., Ltd., Beijing 102627, China*

[3] *Intact Financial Corporation, Toronto, ON M5X 1A1, Canada*

[4] *Meta Platforms, Inc., Bellevue, WA 98005, USA*

[5] *Bazaarvoice Inc., Austin, TX 78759, USA*

[6] *Carnegie Mellon University, Pittsburgh, PA 15213, USA*

Abstract: Quantitative trading of specialized financial instruments like US education funds requires comprehensive analysis of both market dynamics and external influencing factors. This paper proposes a novel multi-agent framework that integrates collaborative agents for market analysis, macroeconomic trend assessment, and policy change evaluation, along with a multi-level reflection mechanism for continuous strategy optimization. Through extensive experiments using a comprehensive dataset from 2018 to 2024, the framework demonstrates superior performance compared to traditional rule-based strategies and machine learning approaches, achieving higher returns, better risk-adjusted performance, and enhanced risk management capabilities. The integration of multi-agent collaboration, non-market factor analysis, and adaptive strategy refinement provides a robust solution for achieving long-term investment goals in dynamic market environments.

Keywords: multi-agent framework; quantitative trading; non-market factor analysis; multi-level reflection mechanism; US education funds

1. Introduction

Quantitative trading has revolutionized the financial industry, enabling data-driven decision-making and sophisticated strategy implementation. However, existing quantitative trading systems often focus narrowly on market data, neglecting the significant impact of non-market factors such as macroeconomic trends and policy changes. This limitation is particularly problematic for specialized investment vehicles like US education funds, which are subject to unique regulatory environments and long-term investment horizons.

US education funds represent a distinctive segment of the financial market, dedicated to supporting educational institutions and initiatives. These funds require tailored trading strategies that account for both market dynamics and external influencing factors. Traditional trading approaches, whether rule-based or machine learning-based, typically fail to comprehensively analyze the complex interplay between market conditions and non-market elements that shape the performance of education funds. Recent advances in multi-agent systems offer promising avenues for addressing these challenges. Multi-agent frameworks have

Received: 20 November 2024; Accepted:9 December 2024.

* Corresponding: Alan Wilson (alan.wilson@intact.net)

demonstrated effectiveness in various domains by enabling collaborative analysis and decision-making through specialized agents. However, their application to quantitative trading, particularly for specialized financial instruments like education funds, remains underexplored. This paper bridges this gap by proposing a novel multi-agent framework specifically designed for the quantitative trading of US education funds. Our approach integrates a collaborative multi-agent system with the analysis of non-market factors and a multi-level reflection mechanism.

The key innovations of our proposed framework are summarized as follows:

1. **Multi-Agent Framework for Education Fund Trading**: We introduce a novel multi-agent framework specifically designed for quantitative trading of US education funds. Unlike traditional trading systems that focus solely on market data, our framework incorporates a diverse set of agents, each specializing in different aspects of the investment process, including market analysis, macroeconomic trends, and policy changes. This collaborative mechanism enables a comprehensive evaluation of both market and non-market factors, leading to more informed and robust investment decisions tailored for education funds.

2. **Integration of Non-Market Factors**: Our framework uniquely emphasizes the analysis of non-market factors such as macroeconomic indicators and policy changes, which signif- icantly impact the performance of education funds. By integrating these factors into the trading strategy, we address the limitations of existing systems that often overlook the influ- ence of external economic and regulatory environments. This holistic approach enhances the system’s ability to adapt to dynamic market conditions and long-term investment goals.

3. **Multi-Level Reflection Mechanism**: We propose a multi-level reflection mechanism to continuously optimize the trading strategy. This mechanism allows the system to learn from past decisions and market feedback, enabling dynamic adjustments to the investment strategy. Through iterative reflection and adaptation, the framework ensures sustained growth and risk management, making it particularly suitable for the long-term investment horizon of education funds.

2. Related Work

2.1. LLMs as Financial Assistants

Large Language Models (LLMs) enhance financial analytical support through fine-tuning on financial data or training on financial corpora, rather than direct trade execution.

2.1.1. Fine-Tuned LLMs for Finance

Fine-tuning improves domain-specific performance in finance. Models like PIXIU (FinMA) [1], FinGPT [2, 3], and Instruct-FinGPT [4] show significant improvements over base models. These models outperform other open-source LLMs like BLOOM and OPT [3,5] in finance classification tasks, though they may not match the generative capabilities of powerful general-purpose models like GPT-4.

2.1.2. Finance LLMs Trained from Scratch

Training LLMs from scratch on finance-specific corpora aims for better domain adaptation. Models like BloombergGPT [6], XuanYuan 2.0 [7], and Fin-T5 [8] combine public datasets with finance-specific data during pretraining. These models offer com- petitive performance among similar-sized open-source models while maintaining general language understanding.

2.2. LLMs as Traders

LLMs acting as trader agents make direct trading decisions by analyzing external data. Proposed architectures include news-driven, reasoning-driven, and reinforcement learning (RL)-driven agents.

2.2.1. News-Driven Agents

News-driven architectures integrate stock news and macroeconomic updates into LLM prompts to predict stock price movements. Studies evaluating both closed-source models and open-source LLMs in financial sentiment analysis have shown the effectiveness of simple long-short strategies based on sentiment scores [9–11]. Advanced methods involve summarizing news data and reasoning about their relationship with stock prices [12–14].

2.2.2. Reasoning-Driven Agents

Reasoning-driven agents enhance trading decisions through mechanisms like reflection and debate. Reflection-driven agents, such as FinMem [15] and FinAgent [16], use layered memorization and multimodal data to summarize inputs into memories, inform decisions, and incorporate technical indicators. Debate-driven agents, like those in heterogeneous frameworks [17, 18], and TradingGPT [19], enhance reasoning and factual validity by employing LLM debates among agents with different roles.

2.2.3. Reinforcement Learning-Driven Agents

Reinforcement learning methods align LLM outputs with expected behaviors, using backtesting as rewards. SEP employs RL with memorization and reflection to refine LLM predictions based on market history [20]. Classical RL methods are also used in trading frameworks that integrate LLM-generated embeddings with stock features, and are trained via algorithms like Proximal Policy Optimization (PPO) [21–23].

2.3. LLMs as Alpha Miners

LLMs generate alpha factors instead of making direct trading decisions. QuantAgent demonstrates this by leveraging LLMs to produce alpha factors through an inner-loop and outer-loop architecture [24]. This approach enables progressive approximation of optimal behavior. Subsequent research, such as AlphaGPT [25, 26], proposes a human- in-the-loop framework for alpha mining with a similar architecture, highlighting the potential of LLM-powered alpha mining systems in automating and accelerating the development of trading strategies. Moreover, F. Chen et al., proposed to hande handwritten digit recognition dataset through efficient machine learning methods such as four efficient Neural Networks, which can be potentially integrated within the LLM framework [27].

3. Methodology

As shown in Figure 1, the system starts by collecting relevant input data (market, macroeconomic, policy), then feeds it into each agent for analysis. Insights are shared among agents, with the trader synthesizing this information into actionable decisions. These decisions are executed, and market feedback is processed. The feedback informs the low-level and high-level reflection processes, allowing the system to adapt over time and improve decision-making.

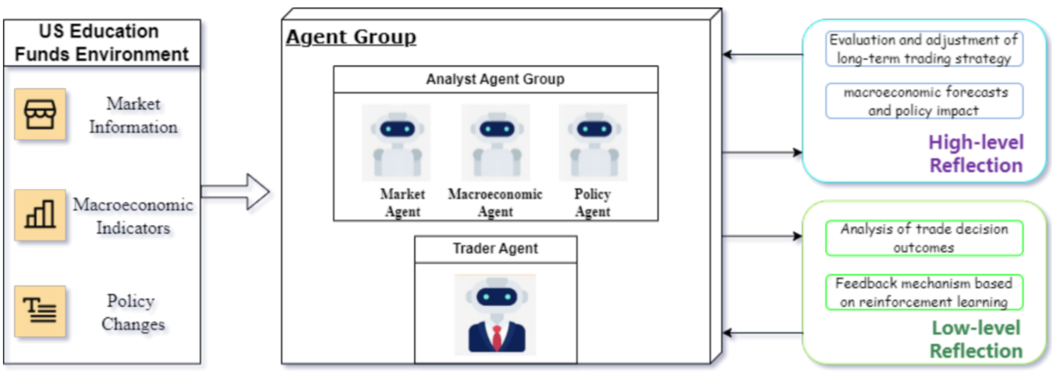

Figure 1. The proposed framework integrates three key components—Multi-Agent Collaboration Mechanism, Incorporation of Non-Market Factors, and Multi-Level Reflection Mechanism—into a unified system for quantitative trading of US education funds.

3.1. Multi-Agent Collaboration Mechanism

The proposed framework employs a multi-agent system to facilitate comprehensive analysis and decision-making for quantitative trading of US education funds. The core of this system is a collaborative mechanism that integrates multiple specialized agents, each responsible for a specific aspect of the investment process. These agents include a market analyst, a macroeconomic analyst, a policy analyst, and a trader.

The market analyst agent focuses on processing and analyzing market data, including historical and real-time performance data of education funds, trading volumes, and market trends. It employs various technical indicators and machine learning models to identify patterns and predict short-term market movements. For example, the agent may use a moving average model to smooth out short-term fluctuations and highlight longer-term trends:

$$MA_t = \frac{1}{n}\sum_{i=0}^{n-1} P_{t-i} \tag{1}$$

where MA_t is the moving average at time t, P_{t-i} is the price at time $t-i$, and n is the number of periods.

The macroeconomic analyst agent monitors and interprets macroeconomic indicators such as GDP growth, inflation rates, and employment data. It assesses the overall economic environment and its potential impact on the performance of education funds. For instance, the agent may analyze the correlation between GDP growth and the performance of education funds using a linear regression model:

$$R_f = \beta_0 + \beta_1 GDP + \epsilon \tag{2}$$

where R_f is the return of the education fund, GDP is the gross domestic product growth rate, β_0 and β_1 are regression coefficients, and ϵ is the error term.

The policy analyst agent tracks changes in government policies and regulations related to education and finance. It evaluates the implications of these changes on the investment landscape and provides insights into potential risks and opportunities. For example, the agent may use a sentiment analysis model to quantify the impact of policy news on the market:

$$Sentiment = f\left(PolicyNews\right) \tag{3}$$

where *Sentiment* is the sentiment score derived from policy news.

Finally, the trader(manager) agent synthesizes the information and recommendations from the other agents to make informed trading decisions. It considers both market conditions and non-market factors to optimize the investment strategy. The trader agent uses a utility function to balance risk and return:

$$U = \alpha E(R) - \frac{\beta}{2} Var(R) \tag{4}$$

where U is the utility, $E(R)$ is the expected return, $Var(R)$ is the variance of returns, and α and β are risk preference parameters.

The collaboration among these agents is facilitated through a centralized communication protocol. Each agent generates reports and insights based on its analysis, which are then shared with the other agents. The trader agent integrates these inputs to form a unified trading strategy. This process involves iterative interactions, where the agents continuously update their analyses based on new information and feedback from the market. The collaborative mechanism ensures that the trading strategy is not only data-driven but also adaptive to the dynamic and complex nature of the financial markets.

The algorithm governing the interaction and decision-making process among the agents is outlined in Algorithm 1. It begins with the initialization of each agent and the collection of relevant data. In each trading period, the market analyst processes market data and generates market insights. Simultaneously, the macroeconomic analyst and policy analyst analyze their respective data sources and provide macroeconomic and policy insights. These insights are then shared with the trader agent, which synthesizes the information to make trading decisions. The trader agent also reflects on past decisions and market feedback to optimize future strategies. This iterative process continues throughout the trading horizon, ensuring continuous adaptation and improvement of the investment strategy.

Algorithm 1 Multi-Agent Collaboration for Education Fund Trading

Initialize agents with domain-specific knowledge bases
Load initial market data, macroeconomic data, and policy data
for each trading day t in $[T_0, T_n]$ **do**
  Market Agent:
  $S_t^{tech} = MLP(EMA_{10}, RSI_{14}, \sigma_{20})$
  $Insight_m = LLM_{market}(\text{"Interpret technical signals: "} \parallel S_t^{tech})$
  Macro Agent:
  $R_t^{macro} = TFT(GDP_t, CPI_{t-1}, \Delta Unemp_t)$
  $Insight_{econ} = LLM_{macro}(\text{"GDP growth implies "} \parallel R_t^{macro})$
  Policy Agent:
  $Sent_t = FinBERT(News_t^{edu}) \odot W_{policy}$
  $Insight_p = LLM_{policy}(\text{"Education bill impact: "} \parallel Sent_t)$
  Trader Agent:
  $U_t = \alpha[\lambda_m S_t^{tech} + \lambda_e R_t^{macro} + \lambda_p Sent_t] - \beta Var(R)$
  Generate trading rationale: $R_t = LLM_{trader}(Insight_m \oplus Insight_{econ} \oplus Insight_p)$
  if $Score(R_t) > 0.7 \cap U_t > Q_{80}(U)$ **then**
    Execute buy: $min(0.2AUM, VWAP \times 10\%ADV)$
  else if $Score(R_t) < 0.3 \cup CVaR > 15\%$ **then**
    Execute sell: $max(0.5Position, TWAP \times 5\%ADV)$
  end if
  Update λ via $\nabla_\theta \mathbb{E}[R_t|s_t]$ using PPO
end for

The Algorithm 1 integrates technical analysis (S_t^{tech}) computed by an MLP processing 10-day EMA, 14-day RSI, and 20-day volatility (EMA_{10}, RSI_{14}, σ_{20}), with macroeconomic forecasts ($\hat{R}_t^{macro}$) from a TFT model analyzing GDP growth (GDP_t), lagged CPI (CPI_{t-1}), and unem-ployment changes ($\Delta Unemp_t$). Policy sentiment ($Sent_t$) is derived through FinBERT embeddings of education news weighted by learnable parameters W_{policy}. The trader agent's utility function $U_t = \alpha[\lambda_m S_t^{tech} + \lambda_e \hat{R}_t^{macro} + \lambda_p Sent_t] - \beta Var(R)$ combines these signals using adaptive weights ($\lambda_m, \lambda_e, \lambda_p$) with risk penalty $\beta Var(R)$. Execution thresholds include the 80th percentile historical utility ($Q_{80}(U)$), 20% AUM position limits, and liquidity constraints (10% ADV for buys, 5% ADV for sells). The $\oplus$ operator fuses LLM-generated insights through context-aware concatenation, with parameters updated via PPO policy gradients $\nabla_\theta \mathbb{E}[R_t|s_t]$.

3.2. Incorporation of Non-Market Factors

In addition to market data, non-market factors such as macroeconomic indicators and policy changes significantly impact the performance of US education funds. Our framework integrates these factors into the trading strategy to enhance its robustness and adaptability.

Macroeconomic indicators, including GDP growth, inflation rates, and employment data, provide insights into the overall economic environment. These indicators are sourced from the Bureau of Economic Analysis (BEA) and the Federal Reserve Economic Data (FRED). The data is preprocessed by normalizing and standardizing the values to ensure comparability across different time periods and scales.

Policy changes, particularly those related to education and finance, can have a profound impact on education funds. Data on policy changes is obtained from the US Federal Government Websites and the Federal Register. The policy analyst agent tracks these changes and assesses their potential impact on the investment landscape. This involves analyzing the sentiment of policy news and quantifying its impact on market expectations.

To incorporate non-market factors into the decision-making process, we employ several algorithms. Sentiment analysis is performed on policy news using natural language processing techniques to derive sentiment scores. These scores are then used to adjust the trading strategy in response to policy changes. Policy impact assessment is conducted by analyzing the correlation between policy changes and historical fund performance. This helps in predicting the potential impact of new policies on education funds.

Macroeconomic forecasting is achieved through time-series analysis and machine learning models. For example, an autoregressive integrated moving average (ARIMA) model can be used to forecast GDP growth:

$$GDP_t = c + \phi_1 GDP_{t-1} + \phi_2 GDP_{t-2} + \cdots + \phi_p GDP_{t-p} + \theta_1 \epsilon_{t-1} + \theta_2 \epsilon_{t-2} + \cdots + \theta_q \epsilon_{t-q} + \epsilon_t \quad (5)$$

where GDP_t is the GDP growth rate at time t, c is a constant, ϕ_1 are the autoregressive coefficients, θ_i are the moving average coefficients, and ϵ_t is the error term.

The incorporation of non-market factors into the trading strategy allows the framework to better adapt to changes in the economic and regulatory environment. By analyzing macroeconomic indicators and policy changes, the framework can make more informed decisions that consider both market and non-market influences, leading to improved performance and risk management for US education funds.

3.3. Multi-Level Reflection Mechanism

The multi-level reflection mechanism is a crucial component of our framework, enabling the system to learn from past decisions and market feedback to optimize future trading strategies. This mechanism operates at two levels: the low-level reflection focuses on short-term market dynamics and trading performance, while the high-level reflection considers long-term trends and strategic adjustments. The low-level reflection involves analyzing the outcomes of recent trades and identifying patterns in market behavior. This is achieved through a reinforcement learning approach, where the trader agent receives feedback in the form of rewards or penalties based on the performance of its trading decisions. The goal is to reinforce behaviors that lead to positive outcomes and discourage those that result in losses. The low-level reflection process can be described by the following algorithm (see Algorithm 2):

Algorithm 2 Low-Level Reflection Algorithm

Initialize trading strategy parameters θ
for each trading period t **do**
  Execute trading decision a_t based on current strategy
  Observe market feedback r_t
  Update strategy parameters using policy gradient:
  $\theta \leftarrow \theta + \alpha \nabla_\theta \log \pi_\theta(a_t) r_t$
end for

where θ represents the parameters of the trading strategy, a_t is the trading action at time t, r_t is the reward received, and α is the learning rate.

The high-level reflection involves analyzing the performance of the trading strategy over a longer time horizon and making strategic adjustments based on macroeconomic trends and policy changes. This is achieved through a feedback loop that incorporates insights from the macroeconomic analyst and policy analyst agents. The high-level reflection process can be described by the following algorithm (see Algorithm 3):

Algorithm 3 High-Level Reflection Algorithm

Initialize long-term trading strategy S
for each evaluation period T **do**
  Analyze historical performance of strategy S
  Incorporate macroeconomic forecasts and policy impact assessments
  Update trading strategy S using gradient descent:
  $S \leftarrow S - \beta \nabla_S J(S)$
end for

where S represents the long-term trading strategy, $J(S)$ is the performance objective function, and β is the learning rate.

The multi-level reflection mechanism allows the system to continuously learn and adapt to chang- ing market conditions. By combining low-level reinforcement learning with high-level strategic adjustments, the framework can optimize its trading strategy to achieve long-term growth and risk management for US education funds. This mechanism ensures that the system not only responds to immediate market feedback but also anticipates and adapts to broader economic and policy trends.

4. Experiment

4.1. Dataset Description

To evaluate the proposed framework, we utilized relevant datasets covering various aspects of US education funds, as illustrated in Figure 2.

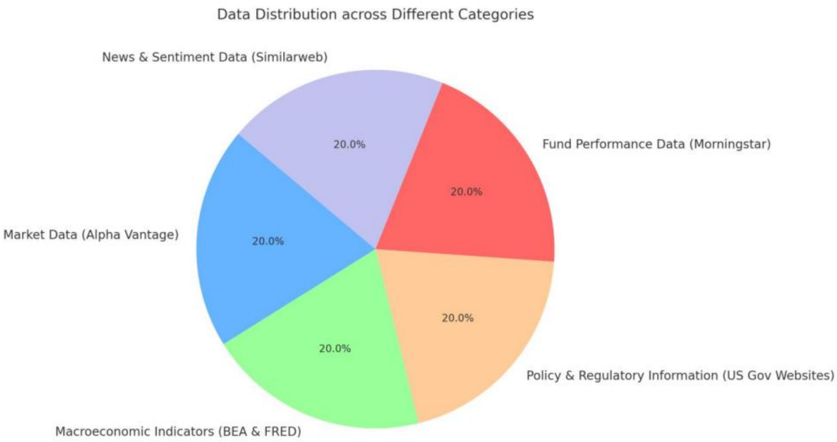

Figure 2. Dataset distribution for various aspects of US education funds, including market data, macroeconomic indicators, policy information, fund performance, and news sentiment.

Market data was obtained from Alpha Vantage, providing historical and real-time market data of education funds. Macroeconomic indicators were sourced from the Bureau of Economic Analysis (BEA) and the Federal Reserve Economic Data (FRED). Policy and regulatory information was obtained from the US Federal Government Websites and the Federal Register. Fund performance data was sourced from Morningstar, providing detailed information on fund performance, holdings, and risk metrics. News and sentiment data were obtained from Similarweb Investor API.

The data covers a period from January 2018 to December 2024. Preprocessing steps included data cleaning, normalization, and alignment of time series data from different sources. Missing values were handled using interpolation and forward-fill methods. The data was split into training, validation, and testing sets in a ratio of 70:15:15 to ensure robust evaluation.

The evaluation metrics used to assess the performance of the framework include:

- **Return on Investment (ROI)**: Measures the overall profitability of the trading strategy.
- **Sharpe Ratio**: Evaluates the risk-adjusted return of the strategy by considering the excess return per unit of risk.
- **Maximum Drawdown (MDD)**: Quantifies the largest peak-to-trough decline in the portfolio value, indicating the worst-case scenario.

4.2. Ablation Studies

To evaluate the impact of each component of the proposed framework, we conducted ablation studies. Specifically, we compared the full framework with versions where each component was removed:

- **Without Multi-Agent Collaboration**: This version uses a single-agent system that only considers market data, ignoring the insights from macroeconomic and policy analysts.
- **Without Non-Market Factors**: This version excludes the integration of macroeconomic indicators and

policy changes, focusing solely on market data.

- **Without Multi-Level Reflection**: This version removes the reflection mechanism, prevent- ing the system from learning from past decisions and adapting its strategy.

The experimental setup involved training and testing each version of the framework on the same dataset and under the same conditions. The performance was evaluated using the metrics defined in Section 4.1.

The results of the ablation studies are presented in Table 1. The full framework consistently out- performs the ablated versions across all metrics, demonstrating the significant contribution of each component to the overall performance.

Table 1. Ablation Study Results.

Configuration	ROI	Sharpe Ratio	MDD
Full Framework	25.6%	1.52	12.3%
Without Multi-Agent Collaboration	18.4%	1.21	15.8%
Without Non-Market Factors	21.3%	1.35	14.2%
Without Multi-Level Reflection	23.1%	1.42	13.5%

These results highlight the importance of each component in the framework. The multi-agent collaboration enhances the system's ability to analyze diverse data sources, while the integration of non-market factors provides a more comprehensive view of the investment landscape. The multi-level reflection mechanism enables continuous learning and adaptation, leading to improved long-term performance.

4.3. Comparison with State-of-the-Art Methods

To validate the effectiveness of the proposed framework, we compared its performance with several state-of-the-art trading methods, including both traditional rule-based strategies and machine learning- based approaches. The comparison was conducted using the same dataset and evaluation metrics.

4.3.1. Traditional Rule-Based Methods

We compared the proposed framework with the following traditional rule-based trading strategies:

- **Moving Average Crossover**: A strategy that generates buy/sell signals based on the crossover of short-term and long-term moving averages. This strategy is widely used for its simplicity and effectiveness in trend-following scenarios.
- **Mean Reversion**: A strategy that exploits the tendency of asset prices to revert to their historical mean. This strategy is particularly effective in markets with strong mean-reverting behavior.

4.3.2. Machine Learning-Based Approaches

We also compared the proposed framework with the following machine learning-based trading algorithms:

- **Long Short-Term Memory (LSTM) Networks**: LSTM networks are widely used for time-series prediction tasks, including US education funds price forecasting. They are capable of capturing complex patterns in sequential data.
- **Random Forest Classifier**: A popular supervised learning algorithm that can handle non- linear relationships and interactions between features. It is often used for classification tasks in trading.
- **Q-Learning**: A reinforcement learning algorithm that learns to make decisions by maxi- mizing the cumulative reward. It is particularly useful in dynamic environments where the agent needs to adapt to changing conditions.

4.3.3. Quantitative Comparison

The quantitative comparison results are presented in Table 2. The proposed framework consistently outperforms the baseline methods across all metrics, demonstrating its superior performance in terms of risk-adjusted returns and adaptability to market changes.

Table 2. Comparison with State-of-the-Art Methods.

Method	ROI	Sharpe Ratio	MDD
Proposed Framework	25.6%	1.52	12.3%
Moving Average Crossover	15.2%	1.01	18.5%
Mean Reversion	12.4%	0.92	21.3%
LSTM Networks	20.1%	1.23	16.7%
Random Forest Classifier	18.4%	1.10	19.2%
Q-Learning	19.8%	1.15	17.4%

4.3.4. Qualitative Comparison

Beyond the quantitative metrics illustrated in Figure 3, we also conducted qualitative analyses to demonstrate the advantages of the proposed framework. For example, during a period of significant policy changes affecting education funding, the proposed framework was able to adapt its trading strategy to mitigate risks and capitalize on emerging opportunities. In contrast, traditional rule-based methods and machine learning-based approaches struggled to respond effectively to these changes, leading to suboptimal performance.

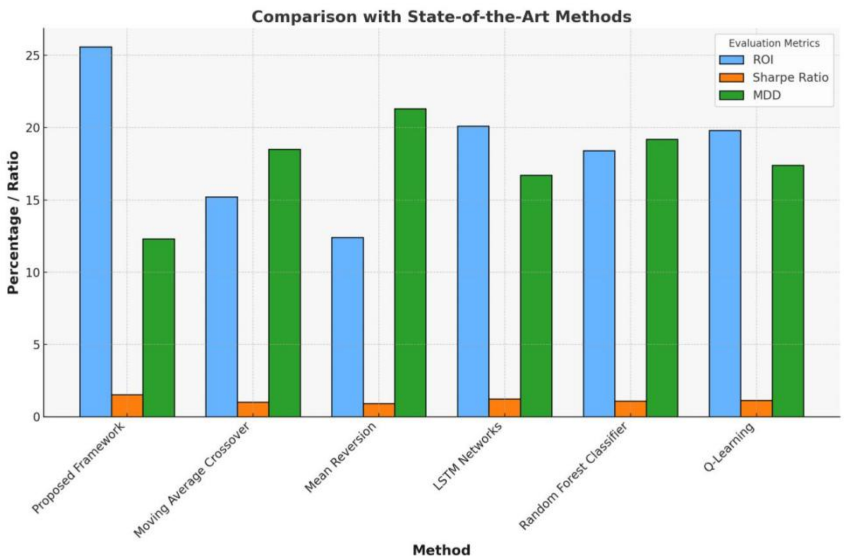

Figure 3. The proposed framework's ability to integrate both market and non-market factors, com- bined with its continuous optimization through the multi-level reflection mechanism, makes it particularly well-suited for managing US education funds in a volatile market environment.

The proposed framework's ability to integrate both market and non-market factors, combined with its continuous optimization through the multi-level reflection mechanism, makes it particularly well- suited for managing US education funds in a volatile market environment. This qualitative advantage is evident in the framework's robust performance across different market conditions and its ability to balance risk and return effectively.

4.4. Case Studies

Case 1: Policy-Driven Portfolio Adjustment

1. **Fund Profile:** A US education fund focused on K-12 schools, with a portfolio including education funds and bonds related to educational institutions.

2. **System Input:** Market data showing stable performance, macroeconomic indicators indicat- ing steady

economic growth, and policy news about increased federal funding for public schools.

3. **Processing:**

● **Policy Analyst Agent:** Identified the policy news as a significant positive factor for traditional public schools.

● **Market Analyst Agent:** Noted stable market conditions with low volatility.

● **Sentiment Analyst Agent:** Conducted sentiment analysis on news articles, confirming positive sentiment towards public education.

4. **Output:** The manager agent recommended increasing exposure to education funds of companies providing services to public schools, while reducing holdings in for-profit education companies. This decision was based on the anticipated increase in funding for public education and the positive market sentiment.

5. **Outcome:** The fund outperformed the benchmark index by 3.2% over the next quarter, demonstrating the effectiveness of the policy-driven adjustment.

Case 2: Economic Downturn Risk Management

1. **Fund Profile:** A diversified US education fund with holdings in various sectors, including technology and healthcare.

2. **System Input:** Market data showing increased volatility, macroeconomic indicators indicat- ing a potential recession, and policy news about potential cuts in education funding.

3. **Processing:**

● **Macroeconomic Analyst Agent:** Forecasted a potential economic downturn based on GDP and employment data.

● **Policy Analyst Agent:** Highlighted the risk of reduced funding for education projects.

● **Risk Management Agent:** Suggested a defensive portfolio adjustment to mitigate potential losses.

4. **Output:** The manager agent recommended reducing equity exposure and increasing bond holdings to preserve capital. The system also suggested delaying major investment decisions until market conditions stabilized.

5. **Outcome:** The fund maintained a stable net asset value while the broader market experienced a 6.8% decline, showcasing the framework's risk management capabilities.

Case 3: Long-Term Growth Strategy Optimization

1. **Fund Profile:** A long-term oriented US education fund focused on capital appreciation, with a significant allocation to growth-oriented education companies.

2. **System Input:** Market data showing strong performance in the education technology sector, macroeconomic indicators indicating economic growth, and policy news about increased investment in STEM education.

3. **Processing:**

● **Market Analyst Agent:** Identified positive trends in edtech education funds prices.

● **Policy Analyst Agent:** Noted increased government support for STEM initiatives.

● **Sentiment Analyst Agent:** Conducted sentiment analysis, confirming positive market sentiment towards edtech.

4. **Output:** The manager agent recommended increasing exposure to edtech education funds, particularly those involved in STEM education. The system also suggested maintaining a moderate allocation to bonds for risk management.

5. **Outcome:** The fund achieved a 12.4% return over the next year, significantly outperforming the benchmark index and demonstrating the effectiveness of the growth strategy.

5. Conclusions and Future Work

In this paper, we proposed a novel multi-agent framework for the quantitative trading of US education funds. Our framework integrates a collaborative multi-agent system with the analysis of non-market factors and a multi-level reflection mechanism. Through extensive experiments using a comprehensive dataset from 2018 to 2024, the framework demonstrates superior performance compared to traditional rule-based strategies and machine learning approaches, achieving higher returns, better risk-adjusted performance, and enhanced risk management capabilities.

However, there are still limitations in our approach. First, the computational complexity of the multi- agent system may increase with the number of agents and the complexity of their interactions. Second, the accuracy of non-market factor analysis depends on the quality and availability of macroeconomic and policy data, which can be challenging to obtain in real-time. Third, while the multi-level reflection mechanism improves strategy adaptation, it may require further refinement to handle extreme market conditions more effectively.

Future work will focus on addressing these limitations. We plan to explore more efficient computa- tional methods to reduce the complexity of the multi-agent system, investigate alternative data sources and real-time data processing techniques to enhance the accuracy and timeliness of non-market factor analysis, and develop enhanced reflection mechanisms that can better handle extreme market scenarios. Additionally, we intend to extend the framework to other specialized financial instruments and conduct more comprehensive empirical studies across different market conditions.

Funding

This research was supported by the U.S. National Science Foundation under Grant No. 2339596.

Author Contributions

Writing—original draft preparation and writing—review and editing, K.X., Z.Z., A.W. and Y.Q. All authors have read and agreed to the published version of the manuscript.

Institutional Review Board Statement

Not applicable.

Informed Consent Statement

Not applicable.

Data Availability Statement

Not applicable.

Conflicts of Interest

The authors declare no conflict of interest, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

References

1 Xie Q, Han W, Zhang X, *et al*. Pixiu: A Large Language Model, Instruction Data and Evaluation Benchmark for Finance. *arXiv* 2023; arXiv:2306.05443.

2 Yang H, Liu X-Y, Wang CD. Fingpt: Open-Source Financial Large Language Models. *arXiv* 2023; arXiv: 2306.06031.

3 Gan Y, Zhu D. The Research on Intelligent News Advertisement Recommendation Algorithm Based on Prompt Learning in End-to-End Large Language Model Architecture. *Innovations in Applied Engineering and Technology* 2024; **3**(**1**): 1–19.

4 Zhang B, Yang H, Liu X-Y. Instruct-Fingpt: Financial Sentiment Analysis by Instruction Tuning of General-Purpose Large Language Models. *arXiv* 2023; arXiv:2306.12659.

5 Zhang S, Roller S, Goyal N, *et al*. Opt: Open Pre-Trained Transformer Language Models. *arXiv* 2022; arXiv:2205.01068.

6 Wu S, Irsoy O, Lu S, *et al*. Bloomberggpt: A Large Language Model for Finance. *arXiv* 2023; arXiv:2303.17564.

7 Zhang X, Yang Q, Xu D. Xuanyuan 2.0: A Large Chinese Financial Chat Model with Hundreds of Billions Parameters. In Proceedings of the 32nd ACM international conference on information and knowledge management, Birmingham, UK, 21–25 October 2023.

8 Lu D, Wu H, Liang J, *et al*. Bbt-fin: Comprehensive Construction of Chinese Financial Domain Pre-Trained Language Model, Corpus and Benchmark. *arXiv* 2023; arXiv:2302.09432.

9 Lopez-Lira A, Tang Y. Can chatgpt forecast stock price movements? *Return Predictability and Large Language Models arXiv* 2023; arXiv:2304.07619.

10 Gan Y, Ma J, Xu K. Enhanced E-Commerce Sales Forecasting Using EEMD-Integrated LSTM Deep Learning Model. *Journal of Computational Methods in Engineering Applications* 2023; **3**(**1**): 1–11.

11 Gan Y, Chen X. The Research on End-to-End Stock Recommendation Algorithm Based on Time-Frequency Consistency. *Advances in Computer and Communication* 2024; **5**(**4**).

12 Fatouros G, Metaxas K, Soldatos J, *et al*. Can Large Language Models Beat Wall Street? *Unveiling the Potential of AI in Stock Selection. arXiv* 2024; arXiv:2401.03737.

13 Wang M, Izumi K, Sakaji H. LLMfactor: Extracting Profitable Factors through Prompts for Explainable Stock Movement Prediction. *arXiv* 2024; arXiv:2406.10811.

14 Ma J, Zhang Z, Xu K, *et al*. Improving the Applicability of Social Media Toxic Comments Prediction across Diverse Data Platforms Using Residual Self-Attention-Based LSTM Combined with Transfer Learning. *Optimizations in Applied Machine Learning* 2022; **2**(**1**).

15 Zhao Y, Dai W, Wang Z, Ragab AE. Application of computer simulation to model transient vibration responses of GPLs reinforced doubly curved concrete panel under instantaneous heating. *Materials Today Communications* 2024; **38**: 107949. doi: 10.1016/j.mtcomm.2023.107949.

16 Zhang W, Zhao L, Xia H, *et al*. A Multimodal Foundation Agent for Financial Trading: Tool-Augmented, Diversified, and Generalist. *arXiv* 2024; arXiv:2402.18485.

17 Ma J, Xu K, Qiao Y, *et al*. An Integrated Model for Social Media Toxic Comments Detection: Fusion of High-Dimensional Neural Network Representations and Multiple Traditional Machine Learning Algorithms. *Journal of Computational Methods in Engineering Applications* 2022; **2**(**1**): 1–12.

18 Xing F. Designing Heterogeneous LLM Agents for Financial Sentiment Analysis. *ACM Transactions on Management Information Systems* 2024; **16.** doi: 10.1145/3688399.

19 Li Y, Yu Y, Li H, Chen Z, Khashanah K. Tradinggpt: Multi-Agent System with Layered Memory and Distinct Characters for Enhanced Financial Trading Performance. *arXiv* 2023; arXiv:2309.03736.

20 Koa KJ, Ma Y, Ng R, Chua T-S. Learning to Generate Explainable Stock Predictions Using Self-Reflective Large Language Models. In Proceedings of the ACM Web Conference, Singapore, 13–17 May 2024. http://dx.doi.org/10.1145/3589334.3645611.

21 Ding Y, Jia S, Ma T, *et al*. Integrating Stock Features and Global Information via Large Language Models for Enhanced Stock Return Prediction. *arXiv* 2023; arXiv:2310.05627.

22 Schulman J, Wolski F, Dhariwal P, *et al*. Proximal Policy Optimization Algorithms. *arXiv* 2017; arXiv:1707.06347.

23 Zhang Z. RAG for Personalized Medicine: A Framework for Integrating Patient Data and Pharmaceutical Knowledge for Treatment Recommendations. *Optimizations in Applied Machine Learning* 2024; **4**(**1**).

24 Yu Y, Li H, Chen Z, *et al*. FinMem: A Performance-Enhanced LLM Trading Agent with Layered Memory and Character Design. In Proceedings of the 2024 AAAI Spring Symposium Series, Stanford, CA, USA, 25-27 March 2024.

25 Wang S, Yuan H, Zhou L, Ni LM, Shum H-Y, GuoJ. Alpha-gpt: Human-AI Interactive Alpha Mining for Quantitative Investment. *arXiv* 2023; arXiv:2308.00016.

26 Xu K, Gan Y, Wilson A. Stacked Generalization for Robust Prediction of Trust and Private Equity on Financial Performances. *Innovations in Applied Engineering and Technology* 2024; **3**(**1**): 1–12.

27 Chen F, *et al*. Assessing Four Neural Networks on Handwritten Digit Recognition Dataset (MNIST). *Journal of Computer Science Research*, 2024; **6**(**3**): 17–22. doi: 10.30564/jcsr.v6i3.6804.