



PDF Download
3701716.3715232.pdf
26 January 2026
Total Citations: 3
Total Downloads: 1311

Latest updates: <https://dl.acm.org/doi/10.1145/3701716.3715232>

RESEARCH-ARTICLE

HedgeAgents: A Balanced-aware Multi-agent Financial Trading System

XIANGYU LI, South China University of Technology, Guangzhou, Guangdong, China

YAWEN ZENG, ByteDance Ltd., Beijing, China

XIAOFEN XING, South China University of Technology, Guangzhou, Guangdong, China

JIN XU, South China University of Technology, Guangzhou, Guangdong, China

XIANGMIN XU, South China University of Technology, Guangzhou, Guangdong, China

Open Access Support provided by:

South China University of Technology

ByteDance Ltd.

Published: 08 May 2025

[Citation in BibTeX format](#)

WWW '25: The ACM Web Conference
2025

April 28 - May 2, 2025
Sydney NSW, Australia

Conference Sponsors:
SIGWEB

HedgeAgents: A Balanced-aware Multi-agent Financial Trading System

Xiangyu Li*
65603605lx@gmail.com
South China University of Technology
Guangzhou, China

Yawen Zeng*
yawenzeng11@gmail.com
ByteDance
Beijing, China

Xiaofen Xing
xfxing@scut.edu.cn
South China University of Technology
Guangzhou, China

Jin Xu
jinxu@scut.edu.cn
South China University of Technology
Pazhou Lab
Guangzhou, China

Xiangmin Xu[†]
xmxu@scut.edu.cn
South China University of Technology
Guangzhou, China

Abstract

As automated trading gains traction in the financial market, algorithmic investment strategies are increasingly prominent. While Large Language Models (LLMs) and Agent-based models exhibit promising potential in real-time market analysis and trading decisions, they still experience a significant -20% loss when confronted with rapid declines or frequent fluctuations, impeding their practical application. Hence, there is an imperative to explore a more robust and resilient framework. This paper introduces an innovative multi-agent system, **HedgeAgents**, aimed at bolstering system robustness via “hedging” strategies. In this well-balanced system, an array of hedging agents has been tailored, where HedgeAgents consist of a central fund manager and multiple hedging experts specializing in various financial asset classes. These agents leverage LLMs’ cognitive capabilities to make decisions and coordinate through three types of conferences. Benefiting from the powerful understanding of LLMs, our HedgeAgents attained a 70% annualized return and a 400% total return over a period of 3 years. Moreover, we have observed with delight that HedgeAgents can even formulate investment experience comparable to those of human experts (<https://hedgeagents.github.io/>).

CCS Concepts

• Information systems → Financial Trading Systems.

Keywords

Quantization Finance, Large Language Model, Multi-agent Systems

*Both authors contributed equally to this paper
[†] Corresponding author

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
WWW Companion '25, Sydney, NSW, Australia
© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-1331-6/2025/04
<https://doi.org/10.1145/3701716.3715232>

ACM Reference Format:

Xiangyu Li, Yawen Zeng, Xiaofen Xing, Jin Xu, and Xiangmin Xu. 2025. HedgeAgents: A Balanced-aware Multi-agent Financial Trading System. In *Companion Proceedings of the ACM Web Conference 2025 (WWW Companion '25)*, April 28-May 2, 2025, Sydney, NSW, Australia. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3701716.3715232>

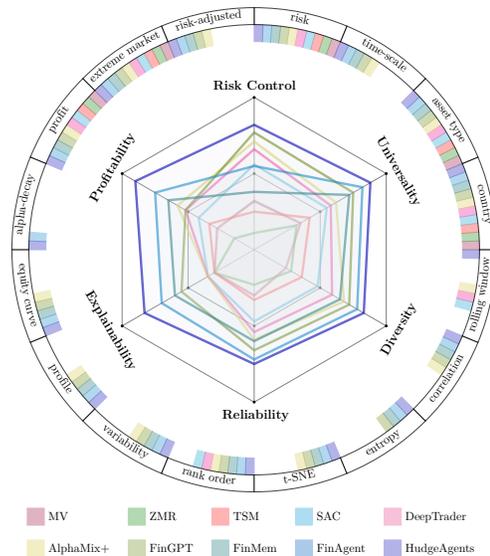


Figure 1: Our method outperformed all baselines on the PRUDEX benchmark [22] across six dimensions. Marks on the inner circle represent the market average, while the outer layer details the measures evaluated.

1 Introduction

In contemporary financial markets, automated trading has emerged as a pivotal investment strategy [3]. Leveraging advanced algorithms, traders can actively track market trends in real-time, enabling swift and precise decision-making, thereby enhancing investment yields while mitigating risk [9]. The integration of artificial intelligence into financial transactions has unlocked significant opportunities, notably with the advent of large language models (LLMs) [4, 26, 36]. These cutting-edge models possess the capability to analyze extensive financial and news information, forecast market variations, and even produce financial analyses

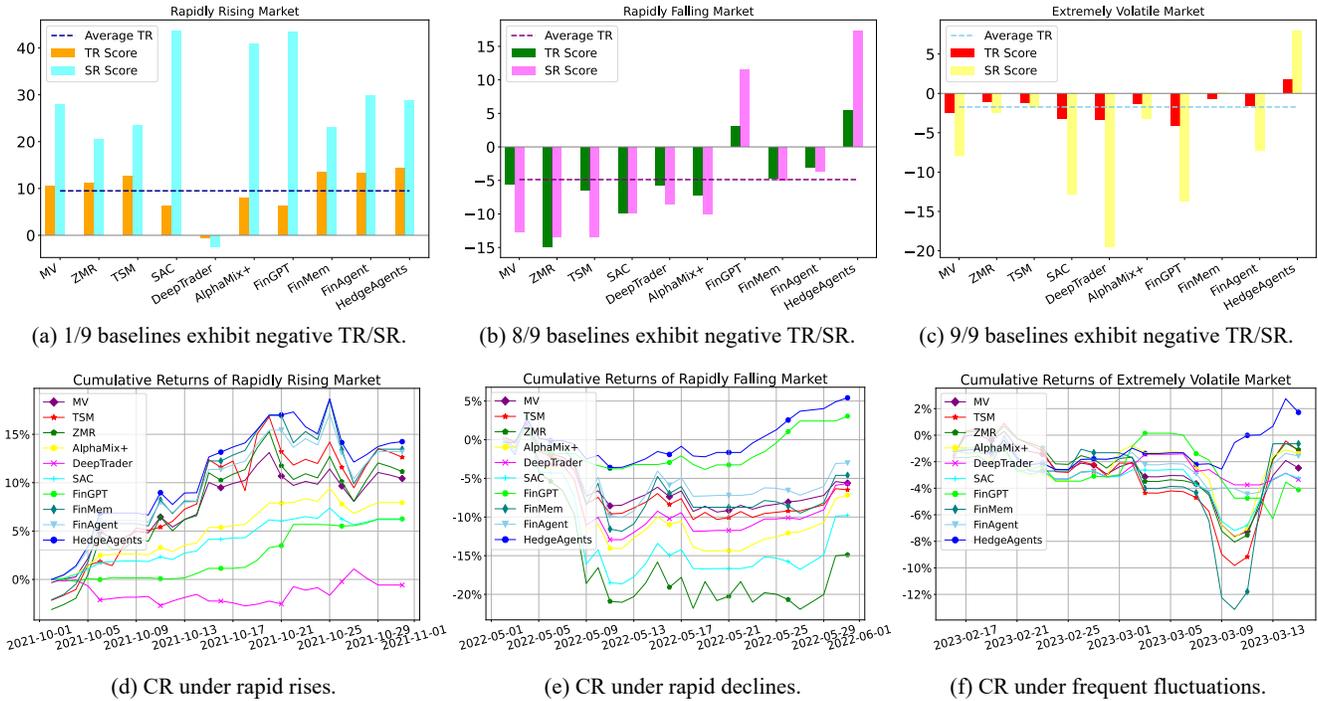


Figure 2: Performance of state-of-the-art models on Total Return (TR), Sharpe Ratio (SR) and Cumulative Returns (CR) in real-world scenarios, with the majority exhibiting negative scores.

and reports [11]. Moreover, LLM-based agents can refine trading strategies through the simulation of market dynamics.

Despite the impressive performance of these cutting-edge models [17, 32], they still lack the robustness required to withstand real-world fluctuations. As depicted in Figure 2, the performance of state-of-the-art models is notably poor on Total Return (TR) and Sharpe Ratio (SR) metrics. For instance, in “rapid decline” scenarios, 8 out of 9 baselines exhibit negative scores below zero. Meanwhile, both RL-based DeepTrader [28] and LLM-based FinGPT [31] experience financial losses ranging from **-15% to -20%** in “frequent fluctuations” scenarios. This limitation stems from the absence of risk management mechanisms within these models or systems, impeding their practical application in real markets. Compounded by the inherent complexity and uncertainty already prevalent in financial markets, the challenge of designing and optimizing trading strategies is further exacerbated.

In this paper, our aim is to incorporate the concept of “hedging” in the development of a robust and reliable financial trading system. Empirically, hedging serves as a strategic tool that enables traders to sustain returns amidst market volatility by constructing diversified portfolios. Along this line, we are committed to building a multi-agent system that mirrors real-world market, with a LLM acting as the brain [7, 30]. A key challenge lies in the configuration and coordination of these hedge agents, which constitutes our primary focus. Illustrated in Figure 2, unconstrained agent system (e.g. FinAgent [38]) can only achieve suboptimal performance.

We are devoted to developing a multi-agent hedging system named **HedgeAgents**. Within this well-balanced framework, we have established a comprehensive array of hedge agents, comprising a fund manager and a team of specialized experts

in various fields such as Stocks, Forex, and Bitcoin. Each expert is tasked with overseeing their designated area, while the fund manager plays a pivotal role in orchestrating their efforts through facilitating discussions, conducting reviews, and consolidating insights. Leveraging the profound understanding capabilities of LLM, these agents (from their memories) have distilled invaluable investment experience akin to those crafted by human counterparts! Drawing from these rich experiences and hedging expertise, our HedgeAgents not only demonstrates unwavering stability even in the face of extreme conditions (Figure 2), but also wins across all metrics, achieving a total return of **400% over 3 years** (Figure 1). We hold a strong belief that this research will pave the way for a dependable automated system that enhances the security of financial investments.

The contributions are summarized as follows:

- To the best of our knowledge, this work represents a pioneering effort in integrating “hedging” within a multi-agent environment to develop a robust and reliable financial trading system.
- We have established a hedging portfolio consisting of three experts and one manager, collaborating through three types of conference, with a LLM serving as the central cognitive hub.
- Extensive experiments has demonstrated that our framework delivers optimal performance on all metrics, achieving a 70% annualized return and a total return of 400% over 3 years. Furthermore, the hedge experts within our framework have produced an invaluable investment experience in their memories akin to human expertise. The codes are available at here¹.

¹<https://hedgeagents.github.io/>

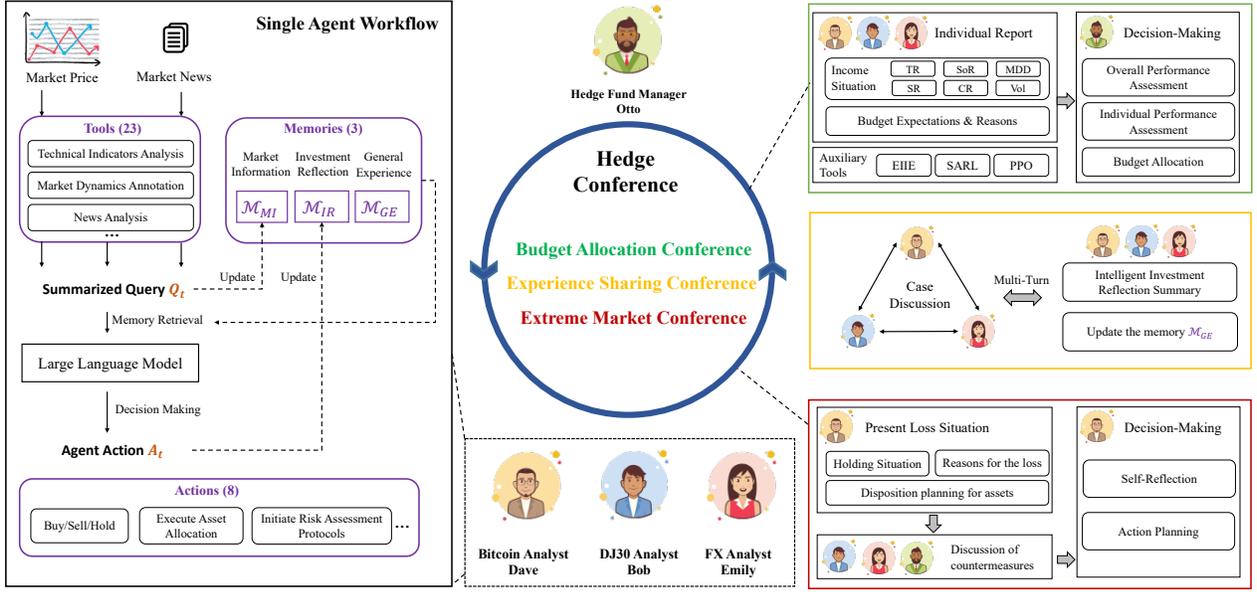


Figure 3: Our HedgeAgents comprise 3 hedging agents and 1 manager. Each agent is equipped with 23 tools and possesses 3 types of memory to execute 8 actions. Furthermore, collaboration among multiple agents can be categorized into three types of conferences: **budget allocation**, **experience sharing**, and **extreme market conference**.

2 Related Work

2.1 Quantitative Finance

Quantitative finance is an interdisciplinary field that merges finance with mathematical and statistical methods to address complex financial challenges [9, 10]. It plays a crucial role in the valuation of sophisticated financial derivatives, portfolio optimization, and the analysis of market dynamics [8, 24]. With the advent of machine learning, the field has seen a surge in predictive modeling capabilities, enhancing both the accuracy of market forecasts and the efficiency of algorithmic trading systems [12], such as MV [33], DeepTrader [28]. As financial markets continue to evolve, the integration of quantitative techniques is essential for developing effective strategies that can withstand market uncertainties [11].

2.2 LLM-based Agent

LLM-based Agents, leveraging the cognitive and generative prowess of Large Language Models [4, 26], promising advancements towards Artificial General Intelligence (AGI)[32]. LLMs play a vital role in enhancing the autonomy[27], reactivity, and social interaction capabilities of agents. This enables agents to perform a range of complex tasks, including natural language interaction, knowledge integration, information memory, logical reasoning, and strategic planning[18, 21]. LLM-based agent systems based on LLMs hold promise for impactful contributions across sectors like finance (e.g. FinGPT [31], FinAgent [38]), offering novel approaches and sophisticated solutions to intricate challenges [7, 30, 37].

3 Proposed Method

3.1 Preliminaries

The objective is to optimize returns while minimizing risk using hedging strategies. Specifically, this task involves utilizing financial

data, encompassing prices and news, as input. Our framework then generates and implements trading actions, such as buying and selling, across three domains: Bitcoin, Stocks, and Forex.

3.2 Overall Framework

The HedgeAgents framework simulates the architecture of a real hedge fund company, aiming to optimize risk hedging for multi-asset investment portfolios. As depicted in Figure 3, the framework comprises four agents: Bitcoin Analyst Dave, Dow Jones Analyst Bob, Forex Analyst Emily, and Hedge Fund Manager Otto. Each of the three analysts is in charge of managing a specific asset. Otto, as their supervisor, is responsible for the overall risk management of the investment portfolio and the distribution of the asset investment budget. To achieve effective collaboration, we have established three types of multi-agent coordination meetings: Budget Allocation Conference (BAC), Experience Sharing Conference (ESC) and Extreme Market Conference (EMC). These conferences serve to facilitate budget allocation, experience summary, and swift emergency actions.

3.3 Definitions of Single Agent

In this section, we will provide a comprehensive overview of the composition and execution process of a single agent, designed to simulate the human decision-making process in investments. Each agent comprises a range of financial analysis tools \mathcal{T} , along with definitions for profile, action \mathcal{A} , memory and reflection \mathcal{M} . Among them, there are 23 tools, encompassing Indicator Analysis, Cryptocurrency Market Analysis, and Risk Management. Actions comprise 8 types, such as Buy/Sell/Hold for current assets. Memory is categorized into 3 types: basic Market Information Memory \mathcal{M}_{MI} , Investment Reflection \mathcal{M}_{IR} , and General Experience \mathcal{M}_{GE} .

<Simplified Prompt Template>

You are {Dave Profile}. The market environment today includes {Prices}, {News}. Through financial analysis tools, {Tool Results} can be obtained. The output format should be JSON, such as {Examples}.

The following will introduce the workflow for a single agent. We have integrated the LLM-based intelligent investment agent into a reinforcement learning framework through a reflection-driven decision-making process, enabling the flexible definition of the reasoning processes. The reflection-driven process includes three steps: memory retrieval, decision making, and reflection update.

Memory Retrieval. We employ retrieval strategy to search for reliable experiences to enhance decision-making. Notably, due to the substantial input received by each agent at time step t (comprising market information, financial news, and tool results), we will first compile a summarized query Q_t and then utilize it to retrieve $K=5$ similar cases \mathcal{M}_{ret} from all memories $\mathcal{M} = \{\mathcal{M}_{MI}, \mathcal{M}_{GE}, \mathcal{M}_{IR}\}$.

Decision Making. Based on the experience \mathcal{M}_{ret} , we extend reinforcement learning as follows:

$$\pi_{\theta^*} = \arg \max_{\pi_{\theta}} \mathbb{E}_{\pi_{\theta}} \left[\sum_{t=0}^T \gamma r_t \mid s_t = s, \mu_t = \mu \right], \quad (4)$$

where r_t is the reward at time step t , which depends on the environment state s_t (e.g. market indicators) and action a_t (e.g. buy/sell). An action a_t is determined through the following process:

$$\pi_{Agent}(a_t \mid s_t, \mu_t) \equiv \mathcal{D}(LLM(\phi_D(s_t, \mathcal{M}_{ret}))), \quad (2)$$

where the prompt template $\phi_D(\cdot)$ is designed to receive responses via a LLM, which are then processed by parsing function $D(\cdot)$ to determine compatible actions in this environment. Therefore, the goal of an agent is to find the policy μ to optimize the total return.

Reflection Update. The basic market information and summarized queries Q_t will be stored in the basic memory \mathcal{M}_{MI} , while the reflection and actions during decision-making will be updated in the investment reflection \mathcal{M}_{IR} .

3.4 Coordination of Hedging Agents

In this section, we design three collaborative meetings to implement the hedging strategies of multiple agents. Among them, the periodic Budget Allocation Conference realizes the trading of multiple assets, with Dave (Bitcoin), Bob (Stocks), Emily (Forex), and manager Otto participating. The Extreme Market Conference serves as an emergency mechanism, enhancing the overall robustness of our system, while the Experience Sharing Conference fosters the exchange of valuable investment insights.

3.4.1 Budget Allocation Conference. This conference spans a 30-day cycle, wherein Dave, Bob, and Emily present reports to Manager Otto aimed at budget allocation. Each report encompasses the current profit situation R_p alongside budget expectations & reasons R_b . Subsequently, manager Otto consolidates all reports and auxiliary tools R_a via a prompt template ϕ_p , to evaluate expected future returns ρ_i .

$$\rho_i = LLM(\phi_p(R_p, R_b, R_a)), i \in \{Dave, Bob, Emily\}. \quad (3)$$

Furthermore, to refine the asset portfolio weighting, three additional indicators are introduced: expected total return I_{etr} , overall portfolio risk I_{pr} , and conditional expected draw down risk I_{cvar} .

$$I_{etr} = \sum_{i=1}^3 \omega_i \rho_i, \quad I_{pr}^2 = \sum_{i=1}^3 \sum_{j=1}^3 \omega_i \omega_j I_{ij}, \quad I_{cvar} = \frac{1}{1-\alpha} \int_1^{\alpha} \text{VaR}(p) dp, \quad (4)$$

where w_i represents the weight of asset i , σ_{ij} is the covariance between the returns of asset i and asset j , and $\text{VaR}(\alpha)$ is the Value at Risk at confidence level α , estimated using the historical simulation method. Finally, budget allocations are determined via an optimization objective function:

$$\omega^* = \arg \max_{\omega} I_{etr} - \lambda_1 I_{pr} - \lambda_2 I_{cvar}, \quad \text{s.t.} \sum_i \omega_i = 1, \omega_i \geq 0, \quad (5)$$

where λ_1 and λ_2 are risk aversion coefficients, which ensure the sum of the allocated budget weights is 1 and non-negative.

3.4.2 Experience Sharing Conference. This is a regular meeting held at the end of each investment cycle, where 3 hedging agents engage in multi-round knowledge accumulation. In each round, each agent will recapitulate a typical case C_t from the memory \mathcal{M}_{IR} for discussion among peers, with all insights D_t being archived in memory \mathcal{M}_{GE} . In this way, the invaluable experience is consolidated to bolster future decision-making.

$$\mathcal{M}_{GE} = LLM(\phi_c(D_t, C_t)), \quad (6)$$

where the LLM aims to update the general experience memory \mathcal{M}_{GE} via a prompt template ϕ_c .

3.4.3 Extreme Market Conference. This is an interim session convened to address volatile market conditions, as identified by manager Otto, characterized by a daily amplitude exceeding 5% or a cumulative three-day amplitude surpassing 10%. During this session, the crisis agent facing challenges are required to present their current portfolio holdings, articulate the reasons behind the crisis, and outline the proposed plans. Following this, Otto and other agents will offer suggestions S_B, S_C, S_E aimed at fostering reflection and optimizing the actions, with the ultimate goal of navigating through the turbulent period effectively.

$$S = (1 - \lambda_3)\phi_B(S_B, S_C) + \lambda_3\phi_B(S_E), \quad (7)$$

where λ_3 is used to balance the importance of suggestions S , and $\phi_B(\cdot)$ is a prompt template. The optimization process is similar to Eqn.2.

$$\pi^* = \arg \max_{\pi(\cdot), \mu(\cdot)} \mathbb{E}_{\pi} \left[\sum_{t=0}^T \gamma r_t \mid s_t = s, \mu_t = \mu \right] \quad (8)$$

s.t. $\pi(a_t \mid s_t, \mu_t) = \mathcal{D}(LLM(\phi_B(s_t, S))) \quad \forall t$

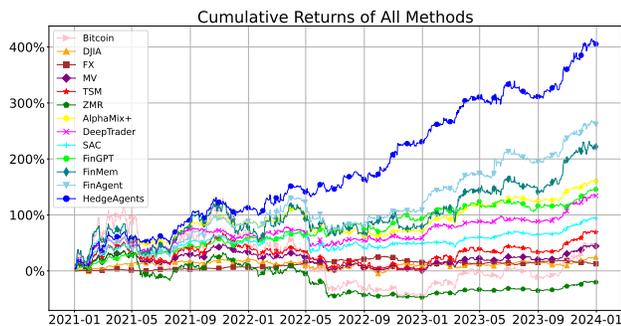
4 Experiments

4.1 Datasets

We have compiled a comprehensive financial dataset comprising Bitcoin, foreign exchange, and the Dow Jones component stocks. These data were sourced from reputable financial databases, namely Yahoo Finance and the Alpaca News API. The dataset spans from

Table 1: Performance comparison of all baselines on 9 evaluation metrics. Bold represents optimal performance, while underline represents suboptimal.

Categories	Models	ARR(%)	TR(%)	SR	CR	SoR	MDD(%)	Vol(%)	ENT	ENB
Market Trends	Bitcoin	12.92	43.97	0.54	1.19	12.49	76.63	3.40	-	-
	FX	4.08	12.74	0.61	0.93	11.99	<u>15.56</u>	0.38	-	-
	DJIA	7.64	24.70	0.59	1.06	11.72	21.94	<u>0.78</u>	-	-
Rule-based	MV	13.03	44.39	0.71	1.25	16.14	32.04	1.13	1.09	1.02
	ZMR	-7.25	-20.21	-0.52	-3.13	-5.15	61.52	1.98	1.55	1.11
	TSM	19.13	69.09	0.78	1.53	18.21	39.14	1.55	1.10	1.09
RL-based	SAC	24.71	93.94	1.16	3.12	23.15	21.56	1.16	1.62	1.14
	DeepTrader	32.78	134.11	1.41	4.06	30.43	20.95	1.21	2.02	1.30
	AlphaMix+	37.59	160.47	1.62	3.69	35.52	25.56	1.17	<u>2.93</u>	1.22
LLM-based	FinGPT	34.22	141.82	<u>1.93</u>	<u>7.64</u>	<u>39.57</u>	17.08	8.76	1.76	1.33
	FinMem	47.67	221.99	1.20	4.02	25.42	32.39	2.16	1.99	1.25
	FinAgent	<u>53.54</u>	<u>261.98</u>	1.80	4.52	39.12	28.24	1.42	2.85	<u>1.41</u>
Ours	HedgeAgents	71.60	405.34	2.41	11.02	58.00	14.21	1.30	3.13	1.53
Improvement(%)		33.75	54.72	24.49	44.28	46.58	16.76	-	6.85	8.51

**Figure 4: Cumulative Returns Comparison of all baselines and our HedgeAgents.**

January 1, 2015, to December 31, 2023, encompassing daily data points such as open, high, low, and close prices, as well as volume and adjusted close prices. Additionally, daily news updates² and 60 standard technical analysis indicators are included for each asset.

4.2 Evaluation Metrics

We compare HedgeAgents and baselines in terms of 9 financial metrics following [19, 23], which include 2 profit metrics: Total Return (TR), Annual Return Rate (ARR); 3 risk-adjusted profit metrics: Sharpe Ratio (SR), Calmar Ratio (CR), Sortino Ratio (SOR); 2 risk metrics: Maximum Drawdown (MDD), Volatility (VOL); and 2 diversity metrics: Entropy (ENT) and Effect Number of Bets (ENB).

²Following [11], we only use processed news headlines that cover the key information of the news.

4.3 Implementation Details

The dataset is divided using January 1, 2021, as the cutoff point, with data from January 1, 2015, to December 31, 2020, designated as the training set, and data from January 1, 2021, to December 31, 2023, as the testing set. Notably, during the testing phase, only historical prices can be accessed to avoid data leakage issues. To ensure the fairness of the evaluation results, all baseline models are trained and tested in the same reinforcement learning environment.

For baseline configurations, we employed Optuna [1] for hyperparameter optimization. For LLM-based approaches and HedgeAgents, we consistently use the GPT-4-1106-preview version with a temperature setting of 0.7 to balance consistency and creativity. The memory module in our framework is designed as a text similarity-based storage and retrieval mechanism, utilizing the text-embedding-3-large model [15], with a top-k value of 5. We ensured LLM-based methods are adjusted to their optimal configurations as suggested in their respective studies.

4.4 Overall Performance Comparison

We selected a diverse range of classical and state-of-the-art baseline models for comparison purposes. These encompass three classical rule-based quantitative investment strategies: MV[33], ZMR[5], and TSM[13]; three reinforcement learning-based financial agents: SAC [6], DeepTrader[29], and AlphaMix+[22]; and three LLM-based methods: FinGPT[31], FinMem[34], and FinAgent[38]. To ensure a fair comparison, we meticulously adhered to the specific requirements of each baseline, providing them with the necessary conditions for optimal performance.

The experimental results are presented in Table 1. The following observations can be made: 1) Compared to rule-based strategies, RL-based methods demonstrate a superior capacity to comprehend

Table 2: Ablation analysis on three conference. ✓denote the inclusion of components.

ESC	BAC	EMC	ARR(%)	TR(%)	SR	CR	SoR	MDD(%)	Vol(%)	ENT	ENB
✓			43.88	197.88	1.90	4.89	42.20	21.70	1.12	2.56	1.23
	✓		45.58	208.52	1.67	4.54	36.36	24.68	1.34	2.39	1.17
		✓	40.97	180.13	1.96	5.07	39.58	19.59	1.02	1.97	1.08
	✓	✓	50.68	242.11	2.01	6.92	45.60	17.26	1.19	2.61	1.28
✓		✓	44.57	202.17	2.24	8.78	55.01	12.02	0.94	2.72	1.39
✓	✓		59.81	308.12	1.93	5.69	40.16	24.44	1.44	2.91	1.42
✓	✓	✓	71.60	405.34	2.41	11.02	58.00	8.68	1.30	3.13	1.53

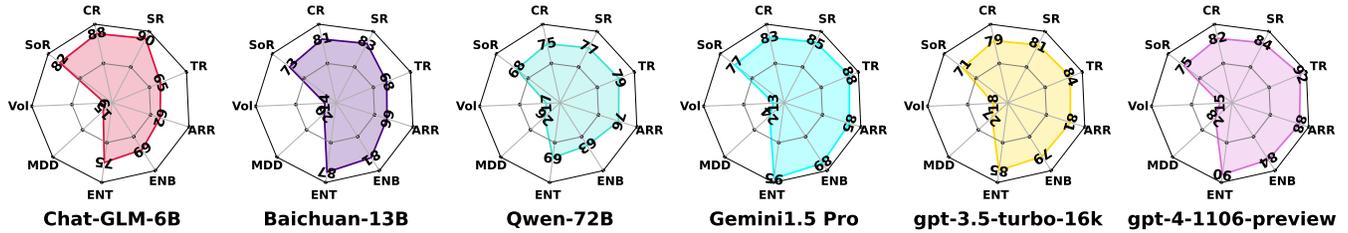


Figure 5: Ablation analysis on several LLM backbones, from open-source to closed-source models.

the intricacies and uncertainties in financial markets. For instance, DeepTrader and AlphaMix+ achieve ARR of 32.78% and 37.59%, respectively, which are significantly higher than rule-based models MV (13.03%), ZMR (-7.25%), and TSM (19.13%), indicating the superior learning and decision-making processes of RL-based agents. 2) Leveraging LLMs, as exemplified by FinGPT, FinMem, and FinAgent, enables the formulation of well-informed decisions, surpassing RL-based methods in risk-adjusted metrics such as SR and CR. FinMem achieves an ARR of 47.67% and a TR of 221.99%, figures that are substantially higher than RL-based agents, suggesting that LLMs’ ability to process and contextualize information leads to more effective investment strategies.

3) Our HedgeAgents exhibits exceptional performance across all metrics, attributable to the following factors: i) **Strategic Budget Allocation.** The dynamic budget allocation plays a pivotal role in optimizing investments across diverse asset categories. This is evidenced by the model achieving the highest ARR of 71.60% and TR of 405.34% among all baselines, showcasing the efficacy of strategic budget distribution in enhancing financial outcomes; ii) **Sophisticated Risk Management.** HedgeAgents demonstrate an impressive MDD of 14.21%, surpassing all other baseline models in risk management. iii) **Diversification and Robustness.** With the highest ENT and ENB, HedgeAgents have demonstrated a diversified investment portfolio that is resilient to the unique risks associated with individual assets. The consistent outperformance of HedgeAgents, as further substantiated by the cumulative returns depicted in Figure 4. Notably, during May 2022, while most portfolios managed by benchmark models experienced significant losses, HedgeAgents adeptly navigated through this challenging period, demonstrating its superior capabilities.

In summary, these analyses confirm that HedgeAgents has not only surpassed benchmarks in individual metrics but also achieved a harmonious balance of returns, risk management, and diversification. This positions HedgeAgents as a cutting-edge solution, adeptly equipped to operate within complex and dynamic investment environments.

4.5 Ablation Study

4.5.1 Effectiveness of Each Conference. To comprehensively evaluate the utility of Experience Sharing Conference (ESC), Budget Allocation Conference (BAC), and Extreme Market Conference (EMC) in our proposed multi-agent architecture, we conducted a systematic ablation study.

The experimental results are shown in Table 2, We have the following observations: 1) The ESC module is crucial for portfolio diversification and hedging tail risks associated with individual assets, which is beneficial for enhancing returns and risk resistance capabilities. When this module is removed, the ENT metric decreases by 16.63% to 2.61, showing the largest drop; although the annualized return and Sharpe ratio decrease, they remain at a moderate level. 2) The BAC module has the greatest impact on returns. After removing this module, the annualized return drops by 37.75% to 44.57%, the lowest among the three single-module experiments; however, when only this module is retained, the annualized return reaches the highest value of 45.58% for a single module. 3) The EMC module is vital for risk aversion and return robustness. After removing this module, the maximum drawdown rate increases by 71.93% to 24.44%, and the Sharpe ratio decreases by 19.86% to 1.93; however, when only this module is retained, both the maximum drawdown rate and Sharpe ratio achieve the best levels among single modules. 4) The synergistic effect of

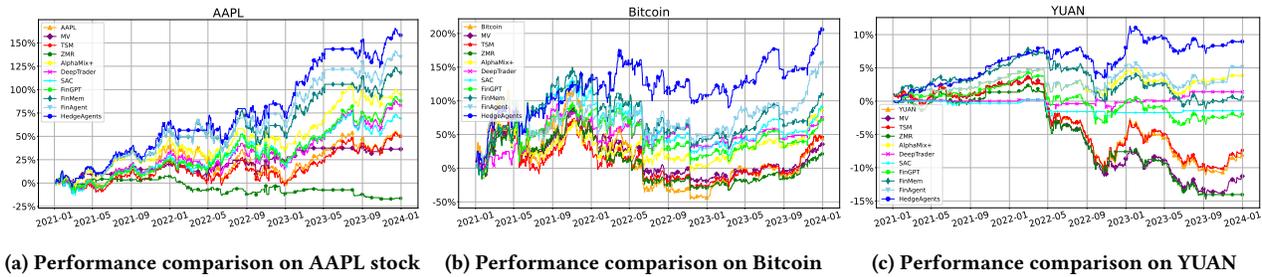


Figure 6: Comparative analysis of cumulative returns for all models in single-asset trading scenarios. These curves show performance differences for AAPL, Bitcoin, and YUAN over a three-year period.

the three modules significantly impacts the overall performance. Taking the Sharpe ratio as an example, compared to the removal of the other modules alone, the Sharpe ratio of the complete architecture improves by 41.29%, 60.65%, and 19.72%, respectively, highlighting the synergistic effect of each module and ensuring a balance between risk and return. This demonstrates the superior performance of our effective multi-agent architecture.

4.5.2 Effectiveness of LLM Backbone. To evaluate the performance of different LLMs as the backbone in our HedgeAgents, we selected 6 representative models, including ChatGLM-6B[35], Baichuan-13B[2], Qwen-72B[20], Gemini 1.5 Pro[25], gpt-3.5-turbo-16k[14], and gpt-4-1106-preview[16]. We incorporated each of these models into the HedgeAgents framework and tested their investment performance.

1) Robustness Analysis: As shown in Figure 5, despite variations in individual metrics across backbones, the curves shapes are strikingly similar, indicating that our framework is not entirely contingent upon the capabilities of a certain language model.

2) Impact of Parameter Size: Experimental results demonstrate that an increase in parameters contributes to enhancing the investment returns. The gpt-4-1106-preview significantly outperforms others in terms of ARR and TR metrics. Notably, models with more parameters tend to adopt more aggressive investment strategies, which can lead to higher risk. Conversely, ChatGLM-6B achieves the best scores in MDD and SR, showcasing strong risk control capabilities.

3) Comparison between Open-source and Closed-source: Closed-source models exhibit excellent performance across all metrics, likely due to their proprietary optimization algorithms and datasets, which provide them with a competitive edge in decision-making capabilities. In the realm of open-source models, Qwen-72B demonstrates performance comparable to the closed-source model.

Therefore, we have selected GPT-4 as the core of our framework. Notably, our system has accumulated a total cost of \$15 over the three years, averaging only 2 cents per day!

4.6 Single-Asset Performance without Hedging

In the previous setup, we examined performance in a multi-asset hedging scenario. In this subsection, we will divide our dataset into three categories: Bitcoin-only, AAPL-only (stocks), and YUAN-only (forex). This will allow us to assess the capabilities of all models under single-asset conditions without any hedging strategy.

Figure 6 illustrates the cumulative returns of all models over a three-year period for each single asset. 1) The Bitcoin-only

setting showcased HedgeAgents' exceptional management of high-volatility assets, achieving a cumulative return of about 210%, far exceeding all baseline models. The Bitcoin expert, Dave, excelled in interpreting blockchain data, anticipating regulatory changes, and analyzing global adoption trends, which facilitated stable growth and resilience during market downturns. 2) In the stock market, represented by AAPL, HedgeAgents attained a cumulative return of approximately 160%, significantly outperforming competitors like FinAgent and FinMem, which recorded returns of about 135% and 120%, respectively. The stock expert, Bob, adeptly navigated both bullish and bearish periods, maintaining consistent growth through his skills in analyzing financial statements and predicting market movements. 3) For the forex market, represented by YUAN, HedgeAgents achieved a modest but positive return of around 9%, surpassing several baselines that showed negative returns. The forex expert, Emily, demonstrated a deep understanding of macroeconomic indicators, geopolitical events, and central bank policies, enabling the model to capitalize on small price movements and achieve profitability.

These results highlight the versatility and robustness of HedgeAgents across various asset classes and market conditions.

4.7 Visualization

To examine our framework's cognitive processes during execution, we have employed visualizations to elucidate the workflow of a single agent and how Otto achieves hedging strategies.

4.7.1 Visualization of single agent. Figure 7 illustrates the decision-making process of a single agent, exemplified by Bitcoin Analyst Dave. Dave's workflow begins with a market analysis noting significant price increases and a robust growth trajectory for Bitcoin. This prompts him to query about historical performance under similar optimistic conditions influenced by macroeconomic policies. Through memory retrieval, Dave recalls past instances where optimistic expectations led to full investment in Bitcoin, but also highlights the challenges of swift position liquidation during subsequent market changes. This historical context informs Dave's decision-making phase, where he considers the positive influence of current macroeconomic policies. Balancing optimism with caution, Dave decides to purchase 50% of the current position in Bitcoin, complemented by continuous market monitoring. This measured approach reflects Dave's synthesis of historical lessons and present market dynamics. The reflection phase reveals a 7.43% increase in asset value over four trading days, validating the effectiveness of Dave's strategy. Consequently, he advocates for maintaining

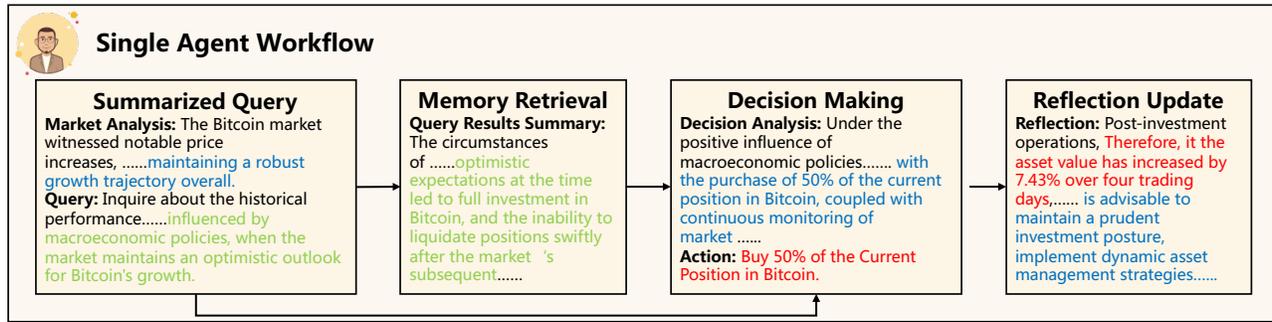


Figure 7: Workflow of a single agent, taking Bitcoin Analyst Dave as an example. Blue represents the analysis content, green indicates the query process, and red shows the specific decisions and outcomes.

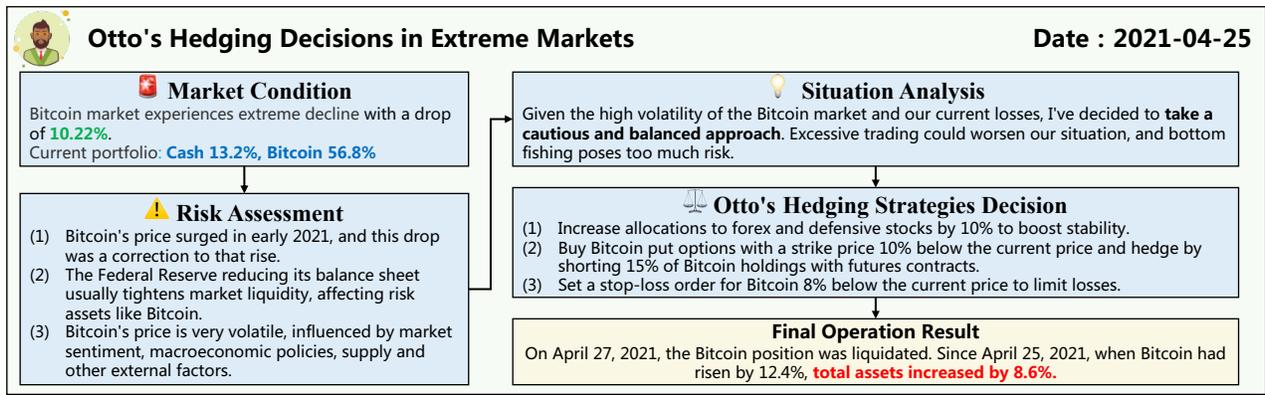


Figure 8: Otto's hedging decisions in extreme markets. Green represents market decline, blue indicates the current portfolio in extreme conditions, and red shows the final execution outcome.

a prudent investment posture and implementing dynamic asset management strategies. Dave's workflow exemplifies an iterative learning process that combines historical insights with real-time analysis, enhancing decision quality in the volatile market.

4.7.2 *Visualization of hedging expert.* Figure 8 illustrates the decision-making process of Otto, our hedging expert agent, during a significant Bitcoin market decline of 10.22%. In this scenario, Otto first assesses the market condition, noting the current portfolio composition of 13.2% cash and 56.8% Bitcoin. Otto's risk assessment identifies three key factors: the recent Bitcoin price surge and subsequent correction, the impact of Federal Reserve policies on market liquidity, and Bitcoin's inherent volatility. Based on this analysis, Otto concludes that a cautious and balanced approach is necessary to protect the portfolio. The situation analysis reveals Otto's concern about the high volatility of the Bitcoin market and the current losses, leading to the decision to avoid excessive trading and risky bottom-fishing strategies. In response to these conditions, Otto proposes a comprehensive hedging strategy. This strategy combines portfolio diversification, options trading, and risk management techniques to create a multi-layered hedge against market volatility. By employing this balanced approach, Otto aims to mitigate downside risk while preserving upside potential in adverse market conditions. The effectiveness of Otto's approach is

evident in the final operation result, where total assets increased by 8.6% over two days, despite the initial market turbulence.

5 Conclusions

This paper presents HedgeAgents, an innovative multi-agent system designed to bolster the robustness of financial actions via hedging strategies. Within this crafted framework, we have developed a suite of specialized hedging agents and organized regular conferences to foster collaboration among multiple agents. Through extensive experiments, our framework has consistently demonstrated superior and robust performance. Furthermore, our observations via visualization indicate that our HedgeAgents is capable of generating investment experience in their memories on par with those achieved by human experts.

6 Acknowledgments

This work is supported in part by the National Natural Science Foundation of China (62372187), in part by the National Key Research and Development Program of China (2022YFC3601005) and in part by the Guangdong Provincial Key Laboratory of Human Digital Twin (2022B1212010004). The authors are highly grateful to the anonymous reviewers for their careful reading and insightful comments.

References

- [1] Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. 2019. Optuna: A Next-Generation Hyperparameter Optimization Framework. In *The 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2623–2631.
- [2] Baichuan Intelligent Technology. 2023. Baichuan-13B. GitHub. Available from: <https://github.com/baichuan-inc/Baichuan-13B/tree/main>.
- [3] Tolga Buz and Gerard de Melo. 2023. Democratization of Retail Trading: Can Reddit's WallStreetBets Outperform Investment Bank Analysts? doi:10.48550/ARXIV.2301.00170
- [4] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. 2023. Palm: Scaling language modeling with pathways. *Journal of Machine Learning Research* 24, 240 (2023), 1–113.
- [5] Louis R. Eeckhoudt and Roger J. A. Laeven. 2018. Dual Moments and Risk Attitudes. arXiv:1612.03347 [q-fin.RM]
- [6] Tuomas Haarnoja, Aurick Zhou, Kristian Hartikainen, George Tucker, Sehoon Ha, Jie Tan, Vikash Kumar, Henry Zhu, Abhishek Gupta, Pieter Abbeel, et al. 2018. Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905* (2018).
- [7] Sirui Hong, Mingchen Zhuge, Jonathan Chen, Xiawu Zheng, Yuheng Cheng, Ceyao Zhang, Jinlin Wang, Zili Wang, Steven Ka Shing Yau, Zijuan Lin, Liyang Zhou, Chenyu Ran, Lingfeng Xiao, Chenglin Wu, and Jürgen Schmidhuber. 2023. MetaGPT: Meta programming for a multi-agent collaborative framework. arXiv:2308.00352 [cs.AI]
- [8] Blanka Horvath, Aitor Muguruza, and Mehdi Tomas. 2021. Deep learning volatility: a deep neural network perspective on pricing and calibration in (rough) volatility models. *Quantitative Finance* 21, 1 (2021), 11–27.
- [9] Takashi Kanamura, Lasse Homann, and Marcel Prokopczuk. 2021. Pricing analysis of wind power derivatives for renewable energy risk management. *Applied Energy* 304 (2021), 117827.
- [10] Gang Kou, Xiangrui Chao, Yi Peng, Fawaz E Alsaadi, Enrique Herrera Viedma, et al. 2019. Machine learning methods for systemic risk analysis in financial sectors. (2019).
- [11] Xiangyu Li, Xinjie Shen, Yawen Zeng, Xiaofen Xing, and Jin Xu. 2024. FinReport: Explainable Stock Earnings Forecasting via News Factor Analyzing Model. In *Companion Proceedings of the ACM Conference 2024* (Singapore, Singapore) (WWW '24). Association for Computing Machinery, New York, NY, USA, 319–327. doi:10.1145/3589335.3648330
- [12] Harald A Mieg. 2022. Volatility as a transmitter of systemic risk: Is there a structural risk in finance? *Risk Analysis* 42, 9 (2022), 1952–1964.
- [13] Tobias J Moskowitz, Yao Hua Ooi, and Lasse Heje Pedersen. 2012. Time series momentum. *Journal of financial economics* 104, 2 (2012), 228–250.
- [14] OpenAI. 2023. GPT-3.5-Turbo: Advanced AI language model by OpenAI. OpenAI API Documentation. <https://platform.openai.com/docs/models/gpt-3-5-turbo>
- [15] OpenAI. 2023. text-embedding-3-large. Available at: <https://openai.com/index/new-embedding-models-and-api-updates/>.
- [16] OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, et al. 2024. GPT-4 Technical Report. arXiv:2303.08774 [cs.CL]
- [17] Keyu Pan and Yawen Zeng. 2023. Do LLMs Possess a Personality? Making the MBTI Test an Amazing Evaluation for Large Language Models. arXiv:2307.16180 [cs.CL]
- [18] Joon Sung Park, Joseph C. O'Brien, Carrie J. Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. 2023. Generative Agents: Interactive simulacra of human behavior. arXiv:2304.03442 [cs.HC]
- [19] Molei Qin, Shuo Sun, Wentao Zhang, Haochong Xia, Xinrun Wang, and Bo An. 2023. Earnhft: Efficient hierarchical reinforcement learning for high frequency trading. *arXiv preprint arXiv:2309.12891* (2023).
- [20] QwenLM. 2023. Qwen-72B. GitHub. Available from: <https://github.com/QwenLM/Qwen/tree/main>.
- [21] Theodore R Sumers, Shunyu Yao, Karthik Narasimhan, and Thomas L Griffiths. 2023. Cognitive architectures for language agents. *arXiv preprint arXiv:2309.02427* (2023).
- [22] Shuo Sun, Molei Qin, Xinrun Wang, and Bo An. 2023. PRUDEX-Compass: Towards Systematic Evaluation of Reinforcement Learning in Financial Markets. arXiv:2302.00586 [q-fin.TR]
- [23] Shuo Sun, Molei Qin, wentao zhang, Haochong Xia, Chuqiao Zong, Jie Ying, Yonggang Xie, Lingxuan Zhao, Xinrun Wang, and Bo An. 2023. TradeMaster: A holistic quantitative trading platform empowered by reinforcement learning. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- [24] Domingo Tavela. 2003. *Quantitative methods in derivatives pricing: an introduction to computational finance*. John Wiley & Sons.
- [25] Gemini Team, Machel Reid, Nikolay Savinov, Denis Teplyashin, et al. 2024. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. arXiv:2403.05530 [cs.CL]
- [26] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971* (2023).
- [27] Lei Wang, Chen Ma, Xueyang Feng, Zeyu Zhang, Hao Yang, Jingsen Zhang, Zhiyuan Chen, Jiakai Tang, Xu Chen, Yankai Lin, Wayne Xin Zhao, Zhewei Wei, and Jirong Wen. 2024. A survey on large language model based autonomous agents. *Frontiers of Computer Science* 18, 6 (March 2024). doi:10.1007/s11704-024-40231-1
- [28] Zhicheng Wang, Biwei Huang, Shikui Tu, Kun Zhang, and Lei Xu. 2021. DeepTrader: a deep reinforcement learning approach for risk-return balanced portfolio management with market conditions Embedding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 643–650.
- [29] Zhicheng Wang, Biwei Huang, Shikui Tu, Kun Zhang, and Lei Xu. 2021. DeepTrader: A Deep Reinforcement Learning Approach for Risk-Return Balanced Portfolio Management with Market Conditions Embedding. *Proceedings of the AAAI Conference on Artificial Intelligence* 35, 1 (May 2021), 643–650. doi:10.1609/aaai.v35i1.16144
- [30] Shijie Wu, Ozan Irsoy, Steven Lu, Vadim Dabravolski, Mark Dredze, Sebastian Gehrmann, Prabhajan Kambadur, David Rosenberg, and Gideon Mann. 2023. Bloomberggpt: A large language model for finance. *arXiv preprint arXiv:2303.17564* (2023).
- [31] Hongyang Yang, Xiao-Yang Liu, and Christina Dan Wang. 2023. FinGPT: Open-Source Financial Large Language Models. *arXiv preprint arXiv:2306.06031* (2023).
- [32] Hui Yang, Sifu Yue, and Yunzhong He. 2023. Auto-GPT for Online Decision Making: Benchmarks and Additional Opinions. arXiv:2306.02224 [cs.AI]
- [33] Jianfeng Yu and Yu Yuan. 2011. Investor sentiment and the mean–variance relation. *Journal of Financial Economics* 100, 2 (2011), 367–381. doi:10.1016/j.jfineco.2010.10.011
- [34] Yangyang Yu, Haohang Li, Zhi Chen, Yuechen Jiang, Yang Li, Denghui Zhang, Rong Liu, Jordan W. Suchow, and Khaldoun Khashanah. 2023. FinMem: A performance-enhanced LLM trading agent with layered memory and character design. arXiv:2311.13743 [q-fin.CP]
- [35] Aohan Zeng, Xiao Liu, Zhengxiao Du, Zihan Wang, Hanyu Lai, Ming Ding, Zhuoyi Yang, Yifan Xu, Wendi Zheng, Xiao Xia, et al. 2022. Glm-130b: An open bilingual pre-trained model. *arXiv preprint arXiv:2210.02414* (2022).
- [36] Yawen Zeng, Da Cao, Xiaochi Wei, Meng Liu, Zhou Zhao, and Zheng Qin. 2021. Multi-Modal Relational Graph for Cross-Modal Video Moment Retrieval. In *Proceedings of the CVPR. IEEE*, 2215–2224.
- [37] Chi Zhang, Zhao Yang, Jiaxuan Liu, Yucheng Han, Xin Chen, Zebiao Huang, Bin Fu, and Gang Yu. 2023. AppAgent: Multimodal agents as smartphone users. arXiv:2312.13771 [cs.CV]
- [38] Wentao Zhang, Lingxuan Zhao, Haochong Xia, Shuo Sun, Jiase Sun, Molei Qin, Xinyi Li, Yuqing Zhao, Yilei Zhao, Xinyu Cai, Longtao Zheng, Xinrun Wang, and Bo An. 2024. A Multimodal Foundation Agent for Financial Trading: Tool-Augmented, Diversified, and Generalist. arXiv:2402.18485 [q-fin.TR]

A Experiment of Ablation Study

A.1 Effectiveness of Each Conference

Cumulative returns of ablation analysis on three conference, as shown in Figure 9 .

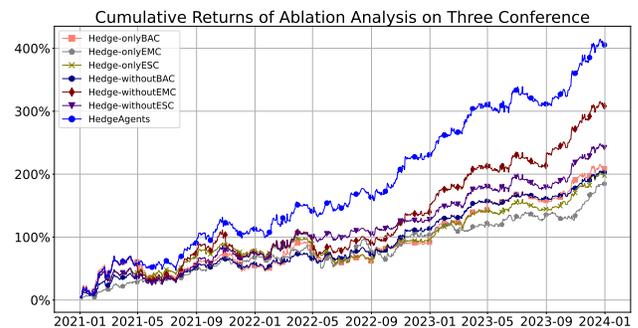


Figure 9: Cumulative Returns of Ablation Analysis on Three Conference

Table 3: Performance comparison of different LLM as the backbone for HedgeAgents on 9 evaluation metrics.

LLM	ARR(%)	TR(%)	SR	CR	SoR	MDD(%)	Vol(%)	ENT	ENB
Chat-GLM-6B	49.14	231.73	2.88	12.26	66.64	9.19	0.78	2.46	1.25
Baichuan-13B	53.24	259.85	2.39	8.85	55.59	13.81	1.02	2.99	1.46
Qwen-72B	61.34	319.99	2.16	8.01	50.21	17.44	1.29	2.35	1.18
Gemini1.5 Pro	68.61	379.37	2.49	11.48	59.94	13.12	1.21	3.31	1.57
gpt-3.5-turbo-16k	65.44	352.81	2.27	8.44	53.49	17.35	1.29	2.93	1.49
gpt-4-1106-preview	71.6	405.34	2.41	11.02	58	14.21	1.3	3.13	1.53

A.2 Effectiveness of LLM Backbone

For using different LLM as the backbone for HedgeAgents, their experimental results are presented in Table 3, and the cumulative returns chart is in Figure 10.

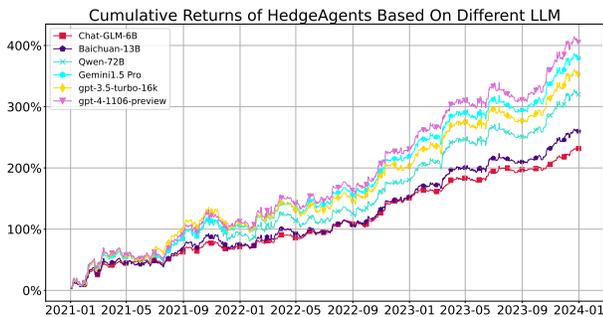


Figure 10: Cumulative Returns of HedgeAgents Based on Different LLM

B Experiment of Memory Retrieval

To quantify the impact of the Memory Retrieval (MR) module on the overall performance of HedgeAgents, we have specifically designed comparative experiments with and without this module.

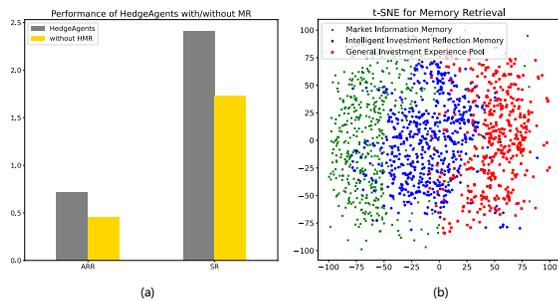


Figure 11: (a) Performance of HedgeAgents with/without MR. (b) Visualization of MR embedding by t-SNE on Dow Jones analyst Bob.

As depicted in Figure 11 (a), the Annualized Return Rate (ARR) and Sharpe Ratio (SR) have both significantly increased upon

the activation of the MR. Compared to the ARR and SR without the module the enhancements reached 57.71% and 39.3%, directly confirming the positive role of the MR module in promoting returns and hedging risks. To delve deeper into the operational mechanics of the MR module, let's consider the case of DJ30 analyst Bob. We'll illustrate the distribution of various types of memory embeddings when employing this module, as depicted in Figure 11 (b). By employing t-SNE for dimensionality reduction and visualization, we have observed that embeddings derived by LLM from hierarchical memory contents, including Market Information Memory, General Investment Experience Pool, and Intelligent Investment Reflection Memory, showcase a discernible clustering distribution in semantic features. This indicates a certain level of distinctiveness between different categories. Such findings substantiate the efficacy of our Hierarchical Memory Retrieval approach in effectively capturing the semantic attributes of memory content, thereby furnishing the agent with valuable contextual memory support.

C Experiment of Temporal Isolation

To minimize the risk of inadvertent information leakage, we designed context-rich prompts incorporating real-time conditions, market variables, and asset-specific data. Figure 12 shows the performance of HedgeAgents during Q1-Q3 of 2024, demonstrating its ability to adapt to rapidly changing market with high predictive accuracy and robustness. The model achieved a Total Return of 68.44%, a Sharpe Ratio of 2.1. These results validate the model's robustness and confirm the effectiveness of our information leakage mitigation strategies.

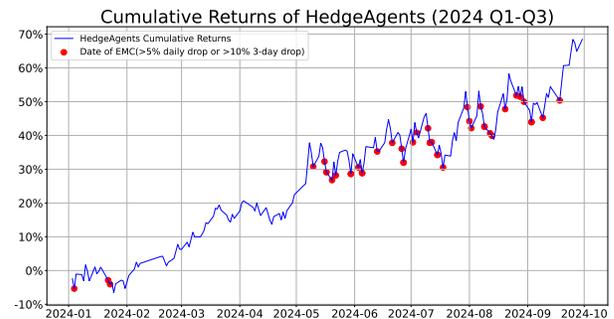


Figure 12: Cumulative Returns of HedgeAgents. Red markers indicate dates of EMC meetings, which occurred 36 times.