

FinHEAR: Human Expertise and Adaptive Risk-Aware Temporal Reasoning for Financial Decision-Making

Jiaxiang Chen^{1*} Mingxi Zou^{1,2*} Zhuo Wang¹ Qifan Wang⁴

Dongning Sun³ Chi Zhang^{2†} Zenglin Xu^{1,5†}

¹Fudan University ²Tensorpacific ³Peng Cheng Laboratory ⁴Meta AI

⁵Shanghai Academy of AI for Science (SAIS)

{jiaxiangchen23, 24210240447}@m.fudan.edu.cn

Abstract

Financial decision-making presents unique challenges for language models, demanding temporal reasoning, adaptive risk assessment, and responsiveness to dynamic events. While large language models (LLMs) show strong general reasoning capabilities, they often fail to capture behavioral patterns central to human financial decisions—such as expert reliance under information asymmetry, loss-averse sensitivity, and feedback-driven temporal adjustment. We propose **FinHEAR**, a multi-agent framework for **Human Expertise** and **Adaptive Risk-aware** reasoning. FinHEAR orchestrates specialized LLM-based agents to analyze historical trends, interpret current events, and retrieve expert-informed precedents within an event-centric pipeline. Grounded in behavioral economics, it incorporates expert-guided retrieval, confidence-adjusted position sizing, and outcome-based refinement to enhance interpretability and robustness. Empirical results on curated financial datasets show that FinHEAR consistently outperforms strong baselines across trend prediction and trading tasks, achieving higher accuracy and better risk-adjusted returns. The code¹ is now publicly accessible.

1 Introduction

Large language models (LLMs) (Achiam et al., 2023; Meta AI, 2024) have demonstrated strong capabilities in general reasoning tasks, including chain-of-thought prompting (Wei et al., 2022), in-context learning (Dong et al., 2024), and code understanding (Nam et al., 2024). However, applying LLMs to financial domains remains challenging due to the dynamic, uncertain, and event-driven nature of real-world markets. Effective financial decision-making requires temporal reasoning, adaptive risk assessment, and responsiveness

*Equal contribution.

†Corresponding author.

¹<https://github.com/pilgrim00/FinHEAR>




Figure 1: Conceptual motivation behind FinHEAR. The framework is grounded in two behavioral finance principles—*information asymmetry* and *prospect theory*—which emphasize the role of expert knowledge and loss-averse risk behavior in financial decision-making.

to external shocks—far beyond the scope of static reasoning benchmarks.

These challenges are further amplified by the complexity of financial data, which is high-frequency, multi-source, and nonlinear (Cao, 2022; Bahoo et al., 2024). Crucially, financial decisions are not purely rational computations—they reflect behavioral patterns shaped by human cognition. Under uncertainty, individuals rely on expert heuristics (Simon, 1955, 1990; Akerlof, 1978), adapting decisions through bounded rationality and accumulated experience. Moreover, risk-sensitive behaviors such as loss aversion and feedback-based

adjustment, as formalized in Prospect Theory (Kahneman and Tversky, 2013; Barberis and Huang, 2001), remain largely unmodeled in existing AI systems. Bridging this gap is essential for aligning LLM reasoning with real-world, high-stakes decision-making.

Prior work has explored rule-based systems, deep learning models, and more recently, LLM-based frameworks for financial decision-making. Rule-based approaches encode expert heuristics but struggle with generalization in dynamic, event-driven environments. Deep learning models, including LSTM (Hochreiter and Schmidhuber, 1997), Transformer (Vaswani et al., 2017), and their variants (Wang and Hao, 2020; Zhou et al., 2021), improve temporal modeling but often lack interpretability and adaptivity. Reinforcement learning (e.g., PPO (Schulman et al., 2017)) offers policy learning under uncertainty, yet remains data-hungry and hard to align with human intuition.

Despite these advancements, existing methods still fall short of key requirements for reliable financial decision-making. First, they lack structured mechanisms to integrate accumulated human expertise—critical for mitigating information asymmetry and supporting decisions under bounded rationality. Second, they exhibit limited sensitivity to risk, often ignoring the loss-averse and risk-adjusted behaviors emphasized by behavioral finance theories such as Prospect Theory. Third, they fail to ensure temporal consistency: unlike static reasoning tasks, financial decisions unfold over time, where early-stage errors can compound if not continually revised. Addressing these gaps is vital for building systems that are not only accurate, but also robust, interpretable, and adaptive to real-world market dynamics.

We introduce **FinHEAR**—a multi-agent framework for **Human Expertise** and **Adaptive Risk-aware** temporal reasoning. FinHEAR addresses three key challenges in financial decision-making: incomplete information, risk sensitivity, and temporal consistency. To reduce information asymmetry, one agent retrieves relevant expert cases to emulate human-guided decisions under bounded rationality—mirroring how investors rely on heuristics under cognitive constraints. To enhance risk awareness, another agent estimates market risk via inter-agent disagreement and adjusts position sizes based on prediction confidence, modeling the loss-averse behavior described by Prospect Theory. A feedback mechanism further refines earlier predictions

using realized outcomes, promoting temporal coherence and limiting error propagation. These components are embedded in a structured, event-driven pipeline, enabling FinHEAR to produce reasoning that is both effective and consistent with behavioral economic theories of human decision-making under uncertainty.

We construct an event-driven financial dataset linking multi-stock price movements with macroeconomic events, firm-level news, and expert commentary. Empirical results show that FinHEAR achieves substantial gains in both trend forecasting and decision making, consistently outperforming strong baselines across key metrics such as ACC, MCC, CR, SR, MDD and CalmarR.

Our key contributions are summarized as follows:

- We propose **FinHEAR**, a multi-agent reasoning framework that integrates **Human Expertise** and **Adaptive Risk-aware** temporal reasoning within a structured, event-driven pipeline—the first to explicitly combine multi-agent coordination with behavioral decision principles for financial reasoning.
- We design explicit mechanisms for expert-guided retrieval, risk-aware action modulation, and feedback-driven temporal refinement—grounded in behavioral decision theories including bounded rationality and loss aversion.
- We validate FinHEAR on real-world financial tasks in both trend forecasting and decision making task, consistently outperforming strong baselines across many metrics.

2 Related Work

Deep Learning Methods. Deep learning has been extensively applied to financial forecasting, with models such as LSTM (Hochreiter and Schmidhuber, 1997), GRU (Cho et al., 2014), and ALSTM (Wang and Hao, 2020) capturing sequential dependencies and emphasizing salient time steps. Transformer-based models (Vaswani et al., 2017) improved long-range modeling, while Informer (Zhou et al., 2021) and TimesNet (Wu et al., 2023) enhanced efficiency and temporal variation modeling. Extensions like StockNet (Gupta et al., 2022) integrated sentiment features, and reinforcement learning methods (Schulman et al., 2017; Kaelbling et al., 1996) introduced reward-driven decision-making, further refined by Logic-Q (Li

et al., 2025) with structured priors.

However, these models primarily focus on pattern extraction and lack mechanisms for event understanding and risk-adaptive reasoning—gaps that **FinHEAR** bridges through expert-informed retrieval, dynamic risk modeling, and feedback-based refinement.

LLM-Based Methods. Recent advances in large language models (LLMs) have enabled structured financial reasoning through prompting and memory-augmented techniques. Methods like Chain-of-Thought (Wei et al., 2022; Wang et al., 2022; Zhang et al.), ReAct (Yao et al., 2023b), and their extensions (Yao et al., 2023a; Sun et al., 2024; Bi et al., 2024; Wu et al., 2024; Zhang et al., 2025) promote stepwise reasoning and agent coordination. FinMEM (Yu et al., 2024a) and FinCon (Yu et al., 2024b) incorporate memory and multi-agent structures for risk-aware decision-making, while TradingAgents (Xiao et al., 2024) mimics real-world organizational hierarchies to improve trading strategies.

While these methods introduce reasoning and coordination capabilities, **FinHEAR** advances further by unifying expert-guided retrieval, dynamic risk modeling, and feedback-based refinement in a behavior-aware framework.

3 Our Method

3.1 Human Expertise Curation

Inspired by the success of legendary investors such as Warren Buffett, George Soros, and Peter Lynch, we integrate structured expert reasoning into **FinHEAR**. These investors embody diverse philosophies—from value investing to macro-driven speculation—offering robust principles for decision-making under uncertainty.

We construct a human-aligned knowledge base by curating writings from expert letters, books, and interviews, segmenting them into passages, and extracting structured entries: the referenced financial event, the expert’s rationale, and contextual attributes (e.g., market condition, asset type, outcome). This enables **FinHEAR** to retrieve analogous cases and ground reasoning in interpretable, expert-informed insights.

Structured Knowledge Construction. Using task-specific prompts and the **GPT-4o** API, we transform the curated corpus into structured *query–knowledge* pairs. The resulting repository

bridges unstructured narratives with actionable decision contexts for downstream reasoning tasks.

Verification and Adaptive Refinement. During inference, real-time financial news is converted into conceptual queries, embedded via **text-embedding-3-large**, and matched to expert queries using cosine similarity. Relevant knowledge activates the **Human Expertise Agent** to guide predictions. Successful predictions reinforce the retrieved knowledge, while failures trigger an adaptive refinement loop via the **Temporal Refine Agent**, which updates expert entries using GPT-4o. This iterative process ensures the knowledge base remains reliable and adaptive to evolving market conditions.

3.2 Architecture of FinHEAR

FinHEAR is a multi-agent reasoning framework for financial decision-making, designed to incorporate structured human expertise, adaptive risk modeling, and feedback-driven temporal refinement. As illustrated in Figure 2, it orchestrates six specialized agents, each aligned with a key component of the reasoning process: capturing historical trends, interpreting current events, retrieving expert analogies, estimating risk, generating trading decisions, and refining predictions based on feedback. These components are grounded in behavioral economic principles such as bounded rationality, information asymmetry, and loss aversion. Further theoretical foundations are detailed in Appendix A.1.

Historical Trend Agent. This agent summarizes sequential market patterns across multiple time horizons, capturing price dynamics, volatility shifts, and cyclical behaviors. These patterns establish the historical context for interpreting new signals. Importantly, the trend representations are dynamically updated using realized outcomes, preventing error accumulation from earlier misinterpretations.

Current Event Agent. This agent interprets real-time financial events—including macroeconomic announcements and firm-specific news—and evaluates their directional influence relative to the evolving market context. The analysis considers both short-term impacts and structural long-term implications, ensuring temporal coherence with the historical trend agent.

Human Expertise Agent. To incorporate analogical reasoning from experienced investors, we construct a curated expert knowledge base. Given a current market state x_{current} , the agent retrieves




Figure 2: FinHEAR architecture. The system coordinates agents for historical trend analysis, event interpretation, and expert retrieval. These agents are organized in a temporal pipeline with feedback that adjusts past analyses based on outcomes. Risk and direction are predicted separately and combined to generate final trading actions. In the multi-asset setting, actions are dynamically aggregated based on risk-aware signals to determine final allocation positions.

the most relevant expert case x_{exp} using event context embeddings. Each context is encoded via a transformer-based encoder $\phi(\cdot)$, and similarity is computed as:

$$x_{\text{retrieved}} = \arg \max_{x_{\text{exp}} \in \mathcal{M}} \cos(\phi(x_{\text{current}}), \phi(x_{\text{exp}})) \quad (1)$$

where $\mathcal{M}$ is the expert memory base. The retrieved case informs high-level decision reasoning based on past expert behavior.

Risk Analysis Agent. This agent evaluates disagreement across the historical, event, and expert agents to estimate overall uncertainty. It outputs a discrete risk level—low, medium, or high—used later to adapt trading actions and the detailed rationale behind it.

Decision Agent. The Decision Agent predicts market direction $d \in \text{up, down}$ using integrated historical, event, and expert signals. This guides the downstream module in selecting trading actions (e.g., long, short, hold, or close) based on both the predicted trend and calibrated risk-aware position scores.

Temporal Refinement Agent. This agent strengthens temporal reasoning by refining past analyses in light of realized outcomes. When predictions conflict with actual market behavior, this agent updates historical trends to correct errors and

ensure consistency across time.

3.3 Adaptive Risk-Aware Decision Making

Single-Asset Setting. Following Prospect Theory, we distinguish three *risk contexts* based on inter-agent coordination and estimated market feedback, and sample a sensitivity score ρ to modulate decision strength. Rather than using uniform distributions, we adopt scaled Beta distributions with asymmetric shapes to reflect context-specific risk preferences:

$$\rho \sim \text{ScaledBeta}(\alpha_r, \beta_r; [a_r, b_r]) \quad (2)$$

In low-risk contexts—often corresponding to potential gain regions— ρ follows a left-skewed distribution (e.g., $\text{Beta}(5, 2)$ on $[0.75, 0.9]$), encouraging moderate risk-taking. In contrast, high-risk contexts adopt a right-skewed shape (e.g., $\text{Beta}(2, 5)$ on $[0.1, 0.4]$), amplifying caution. This design mimics behavioral asymmetries observed in decision-making under uncertainty.

The final position score w is computed by combining risk sensitivity, expert reliability, and event similarity:

$$w = F(\rho, \alpha, \gamma) = f(\rho) \cdot h(\alpha) \cdot g(\gamma) \quad (3)$$

Here, $f(\rho)$ modulates decision strength, which amplifies actions under low risk and attenuating them

under high risk, in line with Prospect Theory’s asymmetric treatment of gains and losses. Expert reliability is given by $h(\alpha) = \max(\alpha, \epsilon_\alpha)$, enforcing a minimum trust threshold under weak signals. For event similarity, $g(\gamma) = \sqrt{\max(\gamma, \epsilon_\gamma)}$ expands low values to avoid over-penalizing partially relevant historical cases. Together, these components stabilize decision outputs under uncertainty and improve robustness in noisy environments.

To select an action, we evaluate the predicted direction $d \in \{\text{up}, \text{down}\}$ together with the computed position score w . We introduce two thresholds, δ_{low} and δ_{high} , to regulate decision-making: the lower threshold suppresses weak or ambiguous signals, while the upper threshold prevents overconfident actions under high risk. When w falls below δ_{low} , the agent chooses to hold; if w exceeds δ_{high} , it chooses to close. For intermediate scores, directional context determines the final action—long if the market is trending upward, or short if trending downward:

$$a = \begin{cases} \text{close}, & w > \delta_{\text{high}} \\ \text{hold}, & w < \delta_{\text{low}} \\ \text{long}, & d = \text{up and } \delta_{\text{low}} \leq w \leq \delta_{\text{high}} \\ \text{short}, & d = \text{down and } \delta_{\text{low}} \leq w \leq \delta_{\text{high}} \end{cases} \quad (4)$$

Multi-Asset Setting. For N assets, each stock i computes a confidence score s_i as in Equation 3, then applies a temperature-scaled softmax:

$$w_i = \frac{\exp(s_i/T)}{\sum_{j=1}^N \exp(s_j/T)}. \quad (5)$$

Here the temperature T adjusts allocation sharpness: lower T sharpens focus, higher T promotes diversification. This generalizes risk-aware decision-making to portfolio-level allocation.

4 Experiments

In this section, we empirically evaluate the performance of **FinHEAR** by addressing the following research questions:

- **RQ1: Forecasting Accuracy.** How accurately can FinHEAR predict the directional trend of major financial assets?
- **RQ2: Backtest Performance.** How well does FinHEAR perform in historical market simulations, in terms of return and risk?
- **RQ3: Component Contribution.** What is the individual impact of each module within FinHEAR on the overall trading outcome?

- **RQ4: Asset-Style Adaptability.** How does FinHEAR’s performance differ for specific assets when guided by human knowledge agents with varying investment philosophies?

4.1 Experimental Setup

Datasets. We primarily focus on three stocks from different sectors: Apple Inc. (AAPL), Tesla, Inc. (TSLA), and Exxon Mobil Corporation (XOM), as well as one commodity asset: Gold (XAUUSD). For these assets, we construct a multi-source financial dataset spanning December 2019 to December 2023, comprising daily prices (open, high, low, close, volume), company financial statements, daily public news and important macroeconomic literature and data such as the official minutes of the U.S. Federal Open Market Committee (FOMC) and periodic reports of the Consumer Price Index (CPI). Each data type is routed to specialized analytical agents based on its temporal and structural characteristics, as detailed in Appendix A.3.

To incorporate long-term investment insights, we compile a corpus of “investor wisdom”—including books, interviews, and historical trading records of iconic investors (e.g., Buffett, Soros)—and dispatch asset-specific queries to retrieval agents aligned with each investor’s philosophy. Further details about human expertise base are available in Appendix A.4.2.

Baselines. To ensure a thorough evaluation across diverse methodological paradigms, we adopt three categories of baselines:

1. **Rule-based methods:** including momentum and mean-reversion strategies, adhering to pre-defined financial trading rules.
2. **Deep learning methods:** including ALSTM, Informer, StockNet, DQN and PPO, representing widely used neural approaches for sequential and decision-making tasks.
3. **LLM-based agents:** including general-purpose generative models(GA), chain-of-thought reasoning approaches(CoT), FinMem and FinCon, covering a broad spectrum of large language model applications in finance.

Evaluation Metrics. We evaluate performance under two complementary settings: (1) **Trend Forecasting**, measured by Accuracy (ACC%) and Matthews Correlation Coefficient (MCC), which assess the correctness and robustness of next-day market predictions. (2) **Decision-making**, measured by Cumulative Return (CR%) (Hull, 2012), Sharpe Ratio (SR) (Sharpe, 1994), Maximum

Table 1: Comparison of the forecasting trend task of FinHEAR and the other 9 baselines in Single-Asset during the test period

Category	Method	XAUUSD		AAPL		TSLA		XOM	
		ACC (%) $\uparrow$	MCC $\uparrow$	ACC (%) $\uparrow$	MCC $\uparrow$	ACC (%) $\uparrow$	MCC $\uparrow$	ACC (%) $\uparrow$	MCC $\uparrow$
Rule-based	Momentum	47.58	-0.049	48.58	-0.0365	51.82	0.0257	47.18	-0.0565
	Mean Reversion	47.35	-0.053	45.49	-0.0742	47.76	-0.0372	52.65	0.0521
Deep-learning	ALSTM	54.32	0.1212	54.69	0.1739	54.88	0.1095	52.25	0.0591
	Informer	53.70	0.1576	53.52	0.1526	55.36	0.0523	53.59	0.0836
	StockNet	52.57	-0.0366	53.59	0.0872	55.77	0.0416	51.96	0.0493
	PPO	54.72	0.1346	53.62	0.1815	54.07	-0.0094	56.52	0.1237
LLM-based	CoT	52.03	0.1493	56.13	0.1629	51.83	0.0623	51.27	0.0585
	GA	54.47	0.1551	54.17	0.1082	56.54	0.1412	52.78	0.0624
	FinMem	55.73	0.1423	55.32	0.1327	58.82	0.1328	53.91	0.0430
	FinCon	56.81	0.1408	52.76	0.0709	57.46	0.1194	56.37	0.1428
Ours	FinHEAR	58.97	0.2017	59.08	0.1964	61.56	0.1897	58.19	0.1625

Drawdown (MDD%) (Ang et al., 2006) and Calmar Ratio(Calmar R) (Young, 1991), which quantify profitability, risk-adjusted return, and resilience under adverse conditions. Formal definitions and formulas are provided in the Appendix A.2.

Implementation Details. To ensure fair comparison, all LLM-based methods, including FinHEAR, use GPT-4o-mini as the underlying model. The data was split into a training period (January 2, 2020 - December 29, 2022) and a test period (January 2, 2023 - December 29, 2023). Both Deep Learning and LLM-based methods were trained on the designated training data and evaluated over the specified test period. Detailed information regarding the data split and evaluation periods is provided in Appendix A.4.1. The detailed experimental configurations for Trend Forecasting and Trading Decision-making are provided in Appendix A.4.4 and Appendix A.4.4, respectively. In the main comparative experiments, all models uniformly adopted a fixed 5-day input time window. This setting was maintained consistently throughout subsequent experiments.

4.2 Experimental Results for Trend Forecasting Task

Despite the potential decoupling between predictive accuracy and financial returns (De Prado, 2018), we argue that accurate trend forecasting is fundamental to the success of financial trading strategies, as predictive quality directly determines directional decisions. Hence, our initial focus is on financial trend prediction as the foundation for strategy development. Specifically, to answer RQ1, we evaluate FinHEAR’s performance in predicting next-day price movements of financial assets.

Table 1 reports trend forecasting results across three datasets. FinHEAR consistently achieves the highest Accuracy (ACC) and Matthews Correlation Coefficient (MCC) (Matthews, 1975), demonstrating strong robustness in single-asset prediction. Overall, rule-based methods exhibit unstable performance due to the rigidity of their rules. Deep learning methods demonstrated relatively stable performance during the subsequent test period. In contrast, LLM-based methods exhibited a wider performance range. This is because their predictions rely on complex strategies or reasoning. While this approach enables potential for high returns, it also renders these methods vulnerable to market changes.

FinHEAR addresses these limitations through two key innovations: expert-guided reasoning to reduce information asymmetry, and feedback-driven refinement to support adaptive temporal decision-making. Its asset-specialized agent design further enhances robustness and interpretability across diverse and volatile market conditions.

4.3 Experimental Results for Trading Decision-making Task

Single-Asset Trading Task To evaluate the trading performance of FinHEAR, we compare it against rule-based strategies, deep learning models, and LLM-based agents across multiple dimensions on four financial assets. As shown in Table 2, FinHEAR consistently achieves the highest Cumulative Return and Sharpe Ratio, indicating superior performance in both return generation and stability. Further details of these experimental results are provided in Appendix A.4.5. This advantage holds across diverse asset types, thereby showcasing the

Table 2: Comparison of the overall trading performance of FinHEAR and the other 9 baselines in Single-Asset during the test period

Category	Method	XAUUSD				AAPL				TSLA				XOM			
		CR (%) ↑	SR ↑	MDD (%) ↓	CalmarR ↑	CR (%) ↑	SR ↑	MDD (%) ↓	CalmarR ↑	CR (%) ↑	SR ↑	MDD (%) ↓	CalmarR ↑	CR (%) ↑	SR ↑	MDD (%) ↓	CalmarR ↑
Rule-based	Momentum	-0.84	-0.28	13.46	-0.06	-4.04	-0.31	19.87	-0.20	40.73	0.85	25.58	1.58	-31.32	-1.56	36.57	-0.86
	Mean Reversion	-19.00	-1.80	19.82	-0.96	-12.15	-0.78	19.78	-0.61	-42.70	-0.89	62.07	-0.69	25.09	0.89	12.90	1.94
Deep-learning	ALSTM	10.45	0.53	14.86	0.70	16.77	0.75	21.49	0.78	11.47	1.14	23.71	0.48	-0.78	-0.36	18.45	-0.04
	Informr	5.36	1.33	17.34	0.31	8.54	1.22	17.34	0.49	-5.42	-1.22	25.12	-0.22	-14.14	-2.19	18.84	-0.75
	StockNet	-19.78	-1.73	23.90	-0.83	12.36	0.85	20.83	0.59	8.97	0.42	19.97	0.45	-5.29	-0.33	27.91	-0.19
	PPO	6.82	1.77	15.85	0.43	15.46	1.67	19.35	0.80	-2.82	-0.01	25.11	-0.11	13.10	1.97	25.93	0.51
LLM-based	CoT	13.98	1.01	18.06	0.77	22.07	1.32	25.46	0.87	-14.26	-1.52	46.61	-0.86	-9.83	-0.51	31.18	-0.32
	GA	14.51	1.78	25.37	0.57	13.66	1.32	15.02	0.91	16.98	0.76	33.98	0.50	-15.59	-1.24	33.74	-0.46
	FinMem	18.32	1.52	18.12	1.01	19.24	1.43	31.02	0.62	28.85	1.22	33.84	0.85	-20.15	-1.59	38.59	-0.52
	FinCon	17.29	1.48	19.25	0.90	-8.17	-0.53	32.47	-0.36	33.73	1.47	29.17	1.16	23.95	1.77	31.72	0.76
Ours	FinHEAR	32.45	1.84	18.47	1.14	30.81	1.95	23.88	1.29	58.74	1.62	25.21	2.33	37.16	2.25	29.61	1.25

framework’s generalizability under varying market dynamics.

FinHEAR addresses these limitations through the integration of expert knowledge to guide decisions, the adjustment of position sizes based on risk perception and behavioral traits like loss aversion, and the refinement of past predictions through feedback. These capabilities enable FinHEAR to adapt effectively in dynamic markets and deliver robust, resilient trading performance.

We further conducted a comprehensive backtest incorporating conservative transaction costs (5 basis points per side). As detailed in Appendix A.4.6, our findings remain robust, with FINHEAR consistently outperforming strong baselines even after accounting for realistic trading frictions.

Robustness Across Market Regimes We evaluate FinHEAR on XAUUSD in 2023 across two phases: a risk-aversion uptrend and a rate-hike downtrend. As shown in Table 3, FinHEAR achieves the highest returns and Sharpe ratios in both regimes and uniquely remains profitable under sustained bearish conditions.

Market Condition	Methods	CR	SR
Risk-Aversion Driven Uptrend (Feb 1 – May 4, 2023)	FinCon	14.38	2.30
	PPO	2.72	1.02
	ALSTM	15.56	4.26
	CoT	19.38	3.63
	FinMem	12.82	2.45
	FinHEAR	18.25	4.79
Rate-Hike Expectation Driven Downtrend (May 5 – Oct 5, 2023)	FinCon	-5.58	-0.14
	PPO	-3.83	-0.06
	ALSTM	-10.39	-2.51
	CoT	-10.12	-2.14
	FinMem	-8.98	-2.27
	FinHEAR	7.82	0.47

Table 3: Performance comparison under different market conditions. CR = cumulative return, SR = Sharpe ratio.

Multi-Asset Portfolio Management Task To evaluate FinHEAR’s capabilities in portfolio management, we extend its multi-agent structure (previously designed for single-asset trading) into a higher-level decision framework capable of handling multiple financial assets simultaneously. This extension enables FinHEAR to coordinate asset-specific reasoning while optimizing overall portfolio performance. We compare FinHEAR with

Table 4: Portfolio Management Performance Evaluation: FinHEAR vs. Baselines (Test Period)

Strategy	CR (%) ↑	SR ↑	MDD (%) ↓	CalmarR ↑
Markowitz	39.21	2.75	12.84	3.05
DQN	45.13	1.60	32.77	1.38
FinCon	55.02	2.11	35.15	1.57
FinHEAR	83.89	3.28	15.77	5.32




Figure 3: Multi-asset portfolio cumulative returns over time for all strategies

three baselines: (1) the classical Markowitz mean-variance (MV) portfolio, where the covariance matrix and expected returns are estimated from training data (Markowitz, 2008); (2) a deep reinforcement learning approach (DQN) (Mnih et al., 2013), trained on data from three years prior to the test period, and employing a continuous action space to reflect real-world trading granularity; and (3) FinCon, a recent LLM-based agent. All methods share the same training and evaluation periods for fairness. As shown in Table 4 and Figure 3, Fin-

Method	XAUUSD	AAPL	TSLA
	CR/SR	CR/SR	CR/SR
FinHEAR (Full, Deep KB)	32.45 / 1.84	30.81 / 1.95	58.74 / 1.62
FinHEAR (Simplified KB)	24.33 / 1.51	19.55 / 1.32	39.26 / 1.51
FinCon	17.29 / 1.48	-8.17 / -0.53	33.73 / 1.47
PPO	6.82 / 1.77	15.46 / 1.67	-2.82 / -0.01
ALSTM	10.45 / 0.53	16.77 / 0.75	11.47 / 1.14
FinMem	18.32 / 1.52	19.24 / 1.43	28.85 / 1.22
CoT (Full KB)	18.19 / 1.45	25.64 / 1.59	10.37 / 0.61
CoT	13.98 / 1.01	22.07 / 1.32	-14.26 / -1.52

Table 5: **Disentangling architecture and knowledge contributions.** Both components are significant, with maximal performance when combined. **CR** = Cumulative Return.

HEAR consistently outperforms all baselines in terms of Cumulative Return (CR), Sharpe Ratio (SR), and Calmar Ratio (CalmarR). These results strongly suggest that FinHEAR’s long-short strategy not only achieves higher excess returns but also better controls extreme risks.

4.4 Ablation Studies

To address RQ3 and RQ4, we conduct two ablation studies under the same training and testing setup as the main experiments. The first study, summarized in Table 6, examines the contribution of key components within FinHEAR. The second study (Table 7) investigates the impact of different expert styles on single-asset trading. Further experimental details are provided in Appendix A.4.7

4.4.1 Ablation Study on Model Components

To assess the role of key components in FinHEAR, we conduct ablation experiments by removing four modules: historical information, human expertise, temporal refinement, and adaptive risk-aware decision-making. Disabling the risk module reduces FinHEAR to a fixed full-position strategy, lacking flexible position control. This study aims to quantify the specific contribution of each module to FinHEAR’s overall performance. As shown in Table 6, all removals cause performance drops, though to varying degrees.

Historical Information and Temporal Refinement. Removing historical information leads to the largest decline in cumulative return and Sharpe ratio, indicating that past news is crucial for capturing market trends. This reflects the idea of *path dependence*—where current behavior is shaped by prior events. Similarly, removing feedback-driven temporal refinement weakens performance by preventing the model from correcting earlier errors,

thereby reducing consistency.

Human Expertise. Removing expert knowledge also degrades performance, especially in single-asset tasks. This supports the concept of *bounded rationality*, where investors rely on expert-informed heuristics under uncertainty. FinHEAR benefits from these inductive biases, improving judgment in volatile or data-limited conditions.

Adaptive Risk-Aware Decision. Without adaptive risk control, FinHEAR becomes less responsive to market uncertainty, resulting in larger drawdowns and lower Calmar Ratios. This supports insights from *prospect theory*, which emphasizes adjusting risk based on potential losses. The risk-aware mechanism helps FinHEAR manage exposure and improve robustness.

4.4.2 Ablation Study on Asset-Expertise Alignment

In our second study, we examine whether different expert investment styles improve single-asset predictions by aligning better with specific asset characteristics. This setup reflects key concepts from behavioral economics, where bounded rationality and information asymmetry lead investors to adopt simplified heuristics based on personal or institutional experience. Table 7 shows clear style-asset alignment such as Buffett-style on AAPL. These results suggest that investment philosophies vary in their effectiveness across different market conditions, and encoding such distinctions can enhance model alignment with asset-specific dynamics.

In addition to investigating expert styles, we also examined how model performance is affected by varying input time window sizes. This study revealed that optimal prediction time windows differ across assets, exhibiting **Temporal Specificity of Asset Performance**. Detailed experimental results and analysis concerning different time windows are provided in Appendix A.4.7.

Overall, these findings underscore FinHEAR’s flexibility in capturing behaviorally grounded strategies through expert-guided specialization. By embedding decision principles such as bounded rationality and loss aversion, FinHEAR supports more interpretable and human-aligned reasoning—positioning it as a behavior-aware framework for expert-informed decision-making in real-world financial contexts.

Table 6: Comparison of trading performance of different variants of FinHEAR under single-asset and portfolio management

Task	Asset	Ablation Variants	CR (%) $\uparrow$	SR $\uparrow$	MDD (%) $\downarrow$	CalmarR $\uparrow$
Single-Asset	AAPL	w/o temporal refinement	22.39	1.67	25.58	0.88
		w/o past trend	14.41	0.86	36.96	0.39
		w/o human expertise	19.75	0.32	35.15	0.56
		w/o adaptive risk-aware	16.32	0.74	31.72	0.51
		FinHEAR	30.81	1.95	23.88	1.29
Portfolio Mgmt.	XAUUSD, AAPL, TSLA, XOM	w/o temporal refinement	52.73	2.17	37.81	1.39
		w/o past trend	37.29	1.34	33.65	1.11
		w/o human expertise	42.12	1.22	28.94	1.46
		w/o adaptive risk-aware	36.18	1.15	35.03	1.03
		FinHEAR	83.89	3.28	25.77	3.22

Table 7: Performance of expert-style agents on single-asset trading

Asset	Expert	CR (%) $\uparrow$	SR $\uparrow$	MDD (%) $\downarrow$	CalmarR $\uparrow$
XAUUSD	Soros	32.45	1.84	28.47	1.34
	Buffet	25.73	1.52	31.20	0.82
	Lynch	28.98	1.63	34.15	0.85
	Graham	21.39	1.33	30.76	0.69
AAPL	Soros	17.35	0.68	34.61	0.50
	Buffet	30.81	1.95	23.88	1.29
	Lynch	26.42	1.57	28.95	0.91
	Graham	28.71	1.98	25.16	1.14
TSLA	Soros	21.26	1.19	38.30	0.55
	Buffet	33.75	1.28	32.11	1.02
	Lynch	48.74	1.62	35.21	1.38
	Graham	35.52	1.49	39.18	1.10
XOM	Soros	25.19	1.54	37.12	0.68
	Buffet	21.62	1.36	25.51	0.85
	Lynch	15.38	0.88	44.01	0.35
	Graham	27.16	2.25	29.61	0.91

4.4.3 Disentangling Architecture vs. Knowledge Base Contributions

We perform a two-path ablation to isolate contributions from architecture and knowledge:

- **FinHEAR (Simplified KB):** Full architecture with a simplified, keyword-only knowledge base.
- **CoT (Full KB):** Standard Chain-of-Thought baseline augmented with our curated knowledge base.

Table 5 shows that: (1) FinHEAR’s architecture improves performance even with minimal knowledge; (2) the curated knowledge base brings substantial independent gains; (3) combining both achieves the best results, confirming their synergy.

4.4.4 Sensitivity to Risk Thresholds

To validate our implementation of Prospect Theory, we conduct a sensitivity analysis on the asymmetric risk thresholds δ_{low} and δ_{high} used in our Risk Agent.

δ_{low}	δ_{high}			
	0.5	0.6	0.7	0.8
0.1	26.88 / 1.71	28.13 / 1.83	29.5 / 1.75	28.05 / 1.48
0.2	29.03 / 1.81	30.81 / 1.95	28.03 / 1.83	30.24 / 1.53
0.3	25.66 / 1.67	25.93 / 1.63	27.55 / 1.92	28.15 / 1.91
0.4	15.34 / 1.23	22.53 / 1.51	24.88 / 1.79	27.89 / 1.99
0.5	–	20.37 / 1.14	21.05 / 1.52	24.11 / 1.36
0.6	–	–	16.97 / 1.39	18.71 / 1.23
0.7	–	–	–	7.36 / 1.24

Table 8: Impact of $(\delta_{\text{low}}, \delta_{\text{high}})$ thresholds on cumulative return (CR) and Sharpe ratio (SR). Each entry shows CR/SR.

As shown in Table 8, the optimal thresholds ($\delta_{\text{low}} = 0.2, \delta_{\text{high}} = 0.6$) align with Prospect Theory’s prediction: agents exhibit stronger risk aversion in loss domains while remaining moderately risk-seeking in gain domains.

5 Conclusion

This paper has proposed FinHEAR to address three key challenges in financial decision-making: incomplete information, risk sensitivity, and temporal consistency. To reduce information asymmetry, it retrieves expert cases to emulate bounded rationality. To manage risk, it adjusts positions based on prediction uncertainty. A feedback mechanism further refines past decisions using outcome-driven updates. These components form a structured, event-driven pipeline grounded in behavioral finance. Experiments on single-asset and portfolio tasks show that FinHEAR consistently outperforms rule-based, deep learning, and LLM-based baselines across forecasting and decision-making metrics. Overall, FinHEAR offers a promising step toward interpretable, resilient, and behavior-aware financial decision-making.

Limitations

While FinHEAR incorporates core principles from behavioral economics, the integration between AI reasoning and human cognitive theories remains preliminary. Future work could explore more intelligent and adaptive selection of expert personas based on asset context, investor profiles, or market conditions, enabling deeper alignment between agent strategies and human decision-making behavior.

Acknowledgments

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, and 1 others. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- George A Akerlof. 1978. The market for “lemons”: Quality uncertainty and the market mechanism. In *Uncertainty in economics*, pages 235–251. Elsevier.
- Andrew Ang, Joseph Chen, and Yuhang Xing. 2006. Downside risk. *The review of financial studies*, 19(4):1191–1239.
- Saqib Bahoo, Marco Cucculelli, Xhemal Goga, Irfan Ullah Hashmi, and Srikanth Manchiraju. 2024. *Artificial intelligence in finance: a comprehensive review through bibliometric and content analysis*. *SN Business & Economics*, 4(1):23.
- Nicholas Barberis and Ming Huang. 2001. Mental accounting, loss aversion, and individual stock returns. *the Journal of Finance*, 56(4):1247–1292.
- Zhenni Bi, Kai Han, Chuanjian Liu, Yehui Tang, and Yunhe Wang. 2024. Forest-of-thought: Scaling test-time compute for enhancing llm reasoning. *arXiv preprint arXiv:2412.09078*.
- Longbing Cao. 2022. Ai in finance: challenges, techniques, and opportunities. *ACM Computing Surveys (CSUR)*, 55(3):1–38.
- Kyunghyun Cho, Bart Van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. 2014. On the properties of neural machine translation: Encoder-decoder approaches. *arXiv preprint arXiv:1409.1259*.
- Marcos Lopez De Prado. 2018. *Advances in financial machine learning*. John Wiley & Sons.
- Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Jingyuan Ma, Rui Li, Heming Xia, Jingjing Xu, Zhiyong Wu, Baobao Chang, and 1 others. 2024. A survey on in-context learning. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 1107–1128.
- Sebastian Ebert and Philipp Strack. 2015. Until the bitter end: On prospect theory in a dynamic context. *American Economic Review*, 105(4):1618–1633.
- Sanford J Grossman and Joseph E Stiglitz. 1980. On the impossibility of informationally efficient markets. *The American economic review*, 70(3):393–408.
- Umang Gupta, Vandana Bhattacharjee, and Partha Sarathi Bishnu. 2022. Stocknet—gru based stock index prediction. *Expert Systems with Applications*, 207:117986.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- John Hull. 2012. *Risk management and financial institutions*,+ *Web Site*, volume 733. John Wiley & Sons.
- Leslie Pack Kaelbling, Michael L Littman, and Andrew W Moore. 1996. Reinforcement learning: A survey. *Journal of artificial intelligence research*, 4:237–285.
- Daniel Kahneman and Amos Tversky. 2013. Prospect theory: An analysis of decision under risk. In *Handbook of the fundamentals of financial decision making: Part I*, pages 99–127. World Scientific.
- Zhiming Li, Junzhe Jiang, Yushi Cao, Aixin Cui, Bozhi Wu, Bo Li, Yang Liu, and Danny Dongning Sun. 2025. Logic-q: Improving deep reinforcement learning-based quantitative trading via program sketch-based tuning. 39(17):18584–18592.
- Harry M Markowitz. 2008. *Portfolio selection: efficient diversification of investments*. Yale university press.
- Brian W Matthews. 1975. Comparison of the predicted and observed secondary structure of t4 phage lysozyme. *Biochimica et Biophysica Acta (BBA)-Protein Structure*, 405(2):442–451.
- Meta AI. 2024. Introducing llama 3.1. <https://ai.meta.com/blog/meta-llama-3-1/>.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. 2013. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*.
- Daye Nam, Andrew Macvean, Vincent Hellendoorn, Bogdan Vasilescu, and Brad Myers. 2024. Using an llm to help with code understanding. In *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering*, pages 1–13.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- William F Sharpe. 1994. The sharpe ratio. *Journal of portfolio management*, 21(1):49–58.

- Hersh Shefrin and Meir Statman. 2000. Behavioral portfolio theory. *Journal of financial and quantitative analysis*, 35(2):127–151.
- Herbert A Simon. 1955. A behavioral model of rational choice. *The quarterly journal of economics*, pages 99–118.
- Herbert A Simon. 1990. Bounded rationality. *Utility and probability*, pages 15–18.
- Linzhuang Sun, Hao Liang, Jingxuan Wei, Bihui Yu, Conghui He, Zenan Zhou, and Wentao Zhang. 2024. Beats: Optimizing llm mathematical capabilities with backverify and adaptive disambiguate based efficient tree search. *arXiv preprint arXiv:2409.17972*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Qi Wang and Yongsheng Hao. 2020. Alstm: An attention-based long short-term memory framework for knowledge base reasoning. *Neurocomputing*, 399:342–351.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2022. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, and 1 others. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.
- Haixu Wu, Tengge Hu, Yong Liu, Hang Zhou, Jianmin Wang, and Mingsheng Long. 2023. Timesnet: Temporal 2d-variation modeling for general time series analysis.
- Jinyang Wu, Mingkuan Feng, Shuai Zhang, Feihu Che, Zengqi Wen, and Jianhua Tao. 2024. Beyond examples: High-level automated reasoning paradigm in in-context learning via mcts. *arXiv preprint arXiv:2411.18478*.
- Yijia Xiao, Edward Sun, Di Luo, and Wei Wang. 2024. Tradingagents: Multi-agents llm financial trading framework. *arXiv preprint arXiv:2412.20138*.
- Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan. 2023a. Tree of thoughts: Deliberate problem solving with large language models. *Advances in neural information processing systems*, 36:11809–11822.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2023b. React: Synergizing reasoning and acting in language models. In *International Conference on Learning Representations (ICLR)*.
- Terry W. Young. 1991. Calmar ratio: A smoother tool. *Futures*. Published by Futures magazine.
- Yangyang Yu, Haohang Li, Zhi Chen, Yuechen Jiang, Yang Li, Denghui Zhang, Rong Liu, Jordan W Su-chow, and Khaldoun Khashanah. 2024a. Finmem: A performance-enhanced llm trading agent with layered memory and character design. In *Proceedings of the AAAI Symposium Series*, volume 3, pages 595–597.
- Yangyang Yu, Zhiyuan Yao, Haohang Li, Zhiyang Deng, Yuechen Jiang, Yupeng Cao, Zhi Chen, Jordan Su-chow, Zhenyu Cui, Rong Liu, and 1 others. 2024b. Fincon: A synthesized llm multi-agent system with conceptual verbal reinforcement for enhanced financial decision making. *Advances in Neural Information Processing Systems*, 37:137010–137045.
- Beichen Zhang, Yuhong Liu, Xiaoyi Dong, Yuhang Zang, Pan Zhang, Haodong Duan, Yuhang Cao, Dahua Lin, and Jiaqi Wang. 2025. Booststep: Boosting mathematical capability of large language models via improved single-step reasoning. *arXiv preprint arXiv:2501.03226*.
- Zhuosheng Zhang, Aston Zhang, Mu Li, and Alex Smola. Automatic chain of thought prompting in large language models. In *The Eleventh International Conference on Learning Representations*.
- Haoyi Zhou, Shanghang Zhang, Jieqi Peng, Shuai Zhang, Jianxin Li, Hui Xiong, and Wancai Zhang. 2021. Informer: Beyond efficient transformer for long sequence time-series forecasting. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pages 11106–11115.

A Appendix

A.1 Behavioral Economics Foundations

FinHEAR is grounded in key principles from behavioral economics, which guide its architectural design and decision-making mechanisms:

- **Bounded Rationality** (Simon, 1955, 1990): Humans have limited cognitive resources and cannot process all available information. FinHEAR incorporates expert heuristics to emulate simplified, experience-driven reasoning processes under uncertainty.
- **Information Asymmetry** (Akerlof, 1978; Grossman and Stiglitz, 1980): Markets often contain uneven distributions of information. By retrieving expert-informed cases and structuring external event signals, FinHEAR reduces decision gaps caused by partial information.
- **Prospect Theory** (Barberis and Huang, 2001; Kahneman and Tversky, 2013; Ebert and Strack, 2015): Human decision-making under risk exhibits loss aversion—losses are weighed more heavily than equivalent gains. FinHEAR models this asymmetry by modulating position sizes according to risk sensitivity and prediction confidence.
- **Behavioral Portfolio Theory** (Shefrin and Statman, 2000): Investors mentally segment their portfolios into layers (e.g., safety vs. growth). FinHEAR reflects this layered decision logic by supporting adaptive strategies across heterogeneous asset classes.

Together, these theoretical foundations support FinHEAR’s goal of aligning algorithmic behavior with real-world investor psychology, promoting interpretability and decision relevance in financial environments.

A.2 Evaluation Metrics

We employ five widely-used metrics in finance to compare the investment rewards of FinHEAR against other algorithmic trading agents. Here are their introductions:

- **Cumulative Return**

Cumulative return (CR) measures a strategy’s profitability over time. In our framework, it is computed by weighting the log return

$\ln(p_{t+1}/p_t)$ —where p_t and p_{t+1} are prices at time t and $t+1$ —by the trading action a_t and position size w_t at each step. Summing over all steps yields the total cumulative return, allowing us to quantify how each decision and exposure contribute to overall performance:

$$CR = \sum_{t=1}^n \ln \left(\frac{p_{t+1}}{p_t} \right) \cdot a_t \cdot w_t \quad (6)$$

- **Sharp Ratio**

The Sharpe Ratio (SR) is a standard metric for evaluating risk-adjusted return. It measures the excess return $R_p - R_f$ relative to its volatility σ_p , where R_p is the average return of the strategy and R_f is the risk-free rate. A higher SR indicates better performance per unit of risk:

$$SR = \frac{R_p - R_f}{\sigma_p} \quad (7)$$

- **Max Drawdown**

The Maximum Drawdown (MDD) quantifies the worst observed loss over a period by measuring the largest peak-to-trough drop in value. Given a peak value P_{peak} and a subsequent trough P_{trough} , MDD reflects the largest percentage decline before recovery:

$$MDD = \max \left(\frac{P_{\text{peak}} - P_{\text{trough}}}{P_{\text{peak}}} \right) \quad (8)$$

- **Calmar Ratio** The Calmar Ratio (CalmarR) evaluates risk-adjusted return using MDD as the risk denominator. It measures how efficiently a strategy generates returns relative to its largest drawdown. Typically, R_p is the annualized strategy return and R_f is the annualized risk-free rate:

$$\text{CalmarR} = \frac{R_p - R_f}{MDD} \quad (9)$$

- **Accuracy**

For evaluating our model’s performance in the forecasting task, we employ Accuracy (ACC). This metric is defined as:

$$ACC = \frac{TP + TN}{TP + TN + FP + FN} \quad (10)$$

True Positives (TP) refer to correctly predicted price increases, while True Negatives (TN) capture correct predictions of no increase. False Positives (FP) arise when an increase

is incorrectly predicted, and False Negatives (FN) indicate missed upward movements. Accuracy (ACC) measures the proportion of correct predictions for the closing price direction.

- **Matthews Correlation Coefficient**

Matthews Correlation Coefficient(MCC) is regarded as a balanced measure for binary classification tasks, particularly useful when class imbalance may be present. The MCC is calculated as follows:

$$MCC = \frac{TP \cdot TN - FP \cdot FN}{\sqrt{(TP + FP)(TP + FN) \cdot (TN + FP)(TN + FN)}} \quad (11)$$

A.3 Dataset Overview

This part provides a detailed description of the curated dataset used in this paper, designed to support financial reasoning and event-aware decision-making. The dataset integrates multi-dimensional information to offer a comprehensive view of the global economic landscape, including macroeconomic indicators, monetary and fiscal policy events, financial market data, industry-specific developments, and expert investment knowledge. These data sources collectively enable structured modeling of market behavior, risk perception, and decision dynamics across multiple temporal and contextual dimensions. The dataset is organized into the following categories:

- **Major Central Bank Communications and Monetary Policy Events:** This category captures policy signals from major central banks worldwide. It includes:
 - Meeting Outcomes (e.g., Federal Funds Rate Decision, ECB Main Refinancing Rate)
 - Post-Meeting Statements (e.g., FOMC Statement, ECB Press Release)
 - Meeting Minutes (e.g., FOMC Minutes, BoE MPC Minutes)
 - Economic Projections (e.g., U.S. SEP, BoJ Outlook Report)
 - Official Speeches and Press Conferences (e.g., Fed Chair Remarks, ECB President Speeches)
- **International Organization and Government Macroeconomic Reports:** Includes

global institutional meetings and national fiscal policy disclosures:

- IMF/World Bank Meetings (e.g., WEO, GFSR)
- G7/G20 Finance Meetings and Communiqués
- National Fiscal Announcements (e.g., U.S. Federal Budget, UK Autumn Statement)

- **Key Macroeconomic Data Releases:** A comprehensive set of high-frequency economic indicators, including:

- Inflation: CPI, PPI, PCE (U.S., Euro Area, China)
- Labor Market: Unemployment Rate, Payrolls, Earnings, Job Openings
- Growth: GDP, Industrial Production, Retail Sales, Durable Goods Orders
- Sentiment: PMI, ISM, Consumer Confidence Indices
- Housing: Housing Starts, Permits, Home Sales
- Trade: Trade Balance (U.S., Euro Area, Germany)

- **Stock Price Time Series (OHLCV):** Daily Open, High, Low, Close, and Volume data for selected equities (e.g., AAPL, TSLA, XOM) from 2020–2023, sourced from Yahoo Finance and Bloomberg.

- **Industry and Company News:** Daily reports (PDF/HTML) on earnings, executive changes, product launches, M&A, and regulatory updates. Sourced from Bloomberg, Reuters, WSJ, FT, etc. Includes 21,329 entries (2020–2023).

- **Current Affairs News:** Daily global news (2020–2023) across politics, military, economy, and society, from major agencies (AP, AFP, BBC, NYT, Xinhua). Includes 22,274 entries.

- **Legendary Investors' Knowledge Base:** Structured corpus of investment philosophy and strategies from prominent investors (e.g., Buffett, Dalio, Soros), including books, interviews, letters, and speeches.

Table 9: Overview of the Dataset

Event Type	Data Type	Source	Frequency	Amount
Central Bank Policy & Comms	PDF/HTML	Fed, ECB, etc.	Various	280
Govt & Intl Macro Reports	PDF/HTML	IMF, G20, etc.	Various	61
Key Macroeconomic Data Releases	PDF/HTML/Text	BLS, U.S. PPI, etc.	Various	2128
Stock daily prices (OHLCV)	CSV	Yahoo Finance, Bloomberg, etc.	Daily	1023
XAUUSD daily price (OHLCV)	CSV	Investing.com, TradingView, etc.	Daily	1035
Industry and Company News	PDF/HTML	Reuters, Bloomberg News, etc.	Daily	21329
Current Affairs	PDF/HTML	AP, Reuters, Bloomberg, etc.	Daily	22274

- **Foundational Books:** Full texts authored by or about legendary investors that codify their investment philosophy, market views, and life experiences. These include:
 - The Intelligent Investor (Benjamin Graham)
 - Security Analysis (Graham & Dodd)
 - One Up on Wall Street, Beating the Street (Peter Lynch)
 - Poor Charlie’s Almanack (Charlie Munger)
 - The Most Important Thing (Howard Marks)
 - The Alchemy of Finance, The New Paradigm for Financial Markets (George Soros)
 - Reminiscences of a Stock Operator (Edwin Lefèvre, based on Jesse Livermore)
 - Common Stocks and Uncommon Profits (Philip Fisher)
 - The Essays of Warren Buffett (compiled by Lawrence Cunningham)
 - Margin of Safety (Seth Klarman)
 - Stocks for the Long Run (Jeremy Siegel)
 - The Dao of Capital (Mark Spitznagel)
 - The Dhandho Investor (Mohnish Pabrai)
- **Shareholder Letters and Memos:** These writings are primary-source reflections of investors’ thinking in real time, often revealing their decision-making logic under changing market conditions:
 - Berkshire Hathaway Shareholder Letters (Warren Buffett, 1977–present)
 - Oaktree Capital Memos (Howard Marks, archived by date and theme)
 - Daily Journal Shareholder Meetings (Charlie Munger, transcript and Q&A)
 - Pershing Square Letters (Bill Ackman)
 - Third Point Letters (Daniel Loeb)
 - Greenlight Capital Letters (David Einhorn)
- **Interviews, Transcripts, and Oral Histories:** These capture dynamic exchanges, spontaneous reasoning, and human intuition in investment contexts:
 - Soros on Soros: Staying Ahead of the Curve (extended Q&A format)
 - Warren Buffett on CNBC/Bloomberg/-Fortune Interviews (indexed by date)
 - Charlie Munger Interviews with Stanford, Caltech, Redlands
 - Peter Lynch on PBS, Fidelity archives
 - Howard Marks Fireside Chats and MasterClass-style interviews
- **Biographies and Third-party Analyses:** These offer synthesized insights into the habits, failures, and competitive edges of key investors, often triangulated with data:
 - The Snowball (Alice Schroeder, biography of Warren Buffett)
 - More Money Than God (Sebastian Mallaby, covers Soros, Jones, Griffin, etc.)
 - The Sages: Buffett, Soros, Volcker and the Maelstrom of Markets (Charles Morris)
 - When Genius Failed (Roger Lowenstein, on LTCM & John Meriwether)
 - The Big Short and Liar’s Poker (Michael Lewis, featuring Paulson, Burry, etc.)
- **Public Speeches and Panel Appearances:** Transcripts and recordings from conferences, universities, or shareholder meetings, often unrehearsed and insight-rich:
 - Berkshire Hathaway Annual Meeting Q&A Archives (1994–present)

- Charlie Munger’s “USC Commencement Speech” and “The Psychology of Human Misjudgment”
 - Peter Thiel at Stanford and YC Startup School (macro-venture angle)
 - Ray Dalio’s “Principles” presentations at TED, Davos, IMF
- **Documentaries and Archival Footage:** Visual content offering non-verbal cues and emotional tone:
 - Becoming Warren Buffett (HBO Documentary)
 - Inside Job (Charles Ferguson)
 - Trader (rare Paul Tudor Jones documentary)
 - The Ascent of Money (Niall Ferguson, includes Soros and others)

These materials form the empirical backbone of a structured Human Expertise Knowledge Base. They enable the modeling of investor archetypes, simulation of investment decision processes, and empirical studies of long-term behavioral consistency and adaptive learning under uncertainty.

A.4 Experimental details

A.4.1 Data Split and Evaluation Periods

To ensure the robustness and effectiveness of our model in a real-world financial environment, we adopted a rigorous time-series split for our training and test sets, explicitly considering a 7-day look-back period requirement for both OHLCV and text data. Our complete dataset spans financial market information from December 26, 2019, to December 29, 2023. This date range ensures that sufficient historical lookback data is available for the earliest predictions.

- **Data Usage During the Training Phase** The model’s actual training and learning, specifically the OHLCV target prediction period, is from January 2, 2020, to December 29, 2022. To make predictions for each trading day within this period, the model incorporates 7 days of preceding OHLCV and text data as historical context. This means that the OHLCV data inputs and text data inputs for the training phase range from December 26, 2019, up to December 29, 2022. Specifically, for any given prediction day within the

training period (e.g., January 2, 2020), the model takes as input the OHLCV and text data from the preceding 7 days (i.e., December 26, 2019, to January 1, 2020) to learn and predict the market dynamics of January 2, 2020. This mechanism ensures that all input information (whether OHLCV or text) available to the model during training is strictly historical relative to the prediction day, rigorously preventing future information leakage.

- **Data Usage During the Testing Phase** After training, the model’s performance is evaluated on an independent and entirely unseen OHLCV period, which runs from January 2, 2023, to December 29, 2023. Consistent with the training phase, to make predictions for each trading day within this testing period, the model also accesses its preceding 7 days of OHLCV and text information. Therefore, the OHLCV data inputs and text data inputs for the testing phase range from December 26, 2022, up to December 29, 2023.

A.4.2 Construct Human Expertise Base

To construct the original expert knowledge base, we first collect textual data associated with the aforementioned domain experts. We then design task-specific prompts to systematically abstract structured query–knowledge pairs from the unstructured corpus, as shown in 6, using Warren Buffett as an example. These prompts are executed via the **GPT-4o** (version: 2024-08-06) API, enabling the extraction of high-level conceptual queries and their corresponding expert-grounded knowledge. This procedure facilitates the transformation of raw expert content into a structured and queryable format suitable for downstream reasoning tasks.

A.4.3 Refine Human Expertise via verification mechanism

To ensure the reliability of the expert knowledge base, we implement a verification mechanism during the testing phase. Specifically, daily news articles are processed to extract real-time conceptual queries, which are then embedded using OpenAI’s **text-embedding-3-large** model. These embeddings are compared against the pre-encoded queries in the knowledge base using cosine similarity. If a sufficiently aligned query is identified—based on a predefined similarity threshold—the corresponding expert knowledge is activated within the **Human Expertise Agent** for prediction. Successful

predictions serve to validate the correctness and applicability of the retrieved expert knowledge. In cases where predictions fail, an adaptive refinement process is triggered. This process leverages feedback information collected by the **Temporal Refine Agent**, which captures contextual cues and decision inconsistencies from the failed prediction. This feedback is then incorporated into a prompt and sent to GPT-4o to generate a revised or augmented version of the expert knowledge entry. Through this iterative refinement loop, the knowledge base is further enriched and aligned with the demands of real-world market conditions, thereby enhancing the original human expert knowledge.

A.4.4 Detailed Configurations in Experiments

- **Trend Forecasting Task Configurations** To ensure a fair comparison for the Trend Forecasting task, all models, including FinHEAR and the various baselines, were configured to output a binary prediction indicating the next-day directional movement of the asset (e.g., "Up" or "Down").

For classification-based baselines, including ALSTM, Informer, and StockNet, their final output layers were designed for two-class prediction. This layer consists of a fully connected layer with two output units followed by a softmax activation, with each class corresponding to either an "Up" or "Down" price movement. At each time step, the model outputs a probability distribution over these two potential outcomes, and the ultimate trend forecast is derived by selecting the class with the highest probability. This standardized two-class output setup ensures consistency in the prediction space across these models.

Similarly, for Reinforcement Learning (RL) based methods such as PPO and DQN, their action spaces were adapted to a binary classification for this task. Instead of continuous position sizing or detailed trading actions, these agents learn to output one of two discrete actions: predicting "Up" or predicting "Down." The reward function in this context is tailored to maximize the accuracy of these directional predictions.

For Large Language Model (LLM) based methods, including the COT agent, GA generative agent, FinMem, and FinCon, their output generation was constrained to produce a

binary forecast. At each decision time step, these agents were prompted or fine-tuned to extract a clear "Up" or "Down" prediction from their text output using a parsing module. This ensures that their output directly corresponds to the two directional classes required for this task.

Through this consistent binary output configuration across all models, we can perform a comprehensive and equitable performance comparison specifically for the trend forecasting task.

- **Trading Task Configurations** To ensure a fair comparison with the FinHEAR model, which is capable of selecting from four distinct trading actions at each decision point—Long, Short, Hold, and Close—we configured the baseline models used in our study accordingly. For the classification-based baselines, including ALSTM, Informer, and StockNet, their final output layers were designed for four-class prediction. This layer consists of a fully connected layer with four output units followed by a softmax activation, with each class corresponding to one of the aforementioned trading actions: 'Long', 'Short', 'Hold', and 'Close'. At each trading time step, the model outputs a probability distribution over these four potential actions, and the ultimate trading signal is typically derived by selecting the action with the highest probability or score. This standardized four-class output setup guarantees consistency in the discrete action space across these models.

Differing in their action space design, we also employed Reinforcement Learning (RL) based methods as baselines capable of more flexible position management. Specifically, we configured agents based on Proximal Policy Optimization (PPO) and Deep Q-Network (DQN). These RL agents learn a trading policy aiming to maximize a cumulative reward signal. The action space configuration varies: PPO utilizes a continuous action space, outputting a target position size within $[-1, 1]$; whereas DQN operates in a discrete action space, with actions encompassing key trading decisions like 'Open Long', 'Open Short', 'Close Position', and 'Hold Position'. The state space for both RL agents comprises rele-

vant market observations (e.g., historical price series, technical indicators) and information about the agent’s current portfolio status (e.g., current position, cash). The reward function guiding the learning process is based on portfolio performance metrics, such as the change in portfolio value or risk-adjusted return over each time step.

As another category of baselines capable of more flexible position management, we included Large Language Model (LLM) based methods, such as the COT agent, GA generative agent, FinMem, and FinCon. These methods leverage LLMs to analyze market data and generate trading recommendations for continuous position management. At each decision time step, these agents accurately extract a concrete, continuous target position size from the text output generated by the LLM using a parsing module. This continuous target position size directly dictates the agent’s trading behavior and intent: a positive target position represents a Long intent; a negative target position indicates a Short intent; and a zero target position corresponds to a Close intent (when holding a position) or maintaining a cash position (Hold) intent (when not holding a position).

Through these diverse baseline configurations spanning different action space granularities, we are able to perform a comprehensive and equitable performance comparison relative to FinHEAR.

A.4.5 Single Asset Trading Result Graphs

as shown in Table 4a, Table 4b, Table 4c, Table 4d, we visualize the cumulative return trajectories for four individual assets—XAUUSD, AAPL, TSLA, and XOM—throughout 2023. These line charts reflect the trading performance using cumulative return as the primary evaluation metric.

A.4.6 Impact of Transaction Costs

To ensure our results are not an artifact of ignoring trading frictions, we conducted a comprehensive backtest under conservative transaction cost assumptions:

1. **New Backtest with Transaction Costs** Following reviewer suggestions, we repeated the AAPL trading experiments, incorporating re-

alistic transaction costs across all methods for a fair comparison.

2. **Conservative Cost Assumption** For every transaction (both entries and exits), we applied a transaction cost of **0.05%** (5 basis points) per side, resulting in a total round-trip cost of **0.1%**. This rate accounts for typical brokerage commissions and average slippage.
3. **Results on AAPL: Performance Advantage Remains Robust** As shown in Table 10, while transaction costs reduce absolute returns across all strategies, FINHEAR continues to significantly outperform all strong baselines in both Cumulative Return (CR) and Sharpe Ratio (SR). This demonstrates that our framework’s ability to generate alpha is not dependent on unrealistic cost assumptions.


Method	No Costs		With Costs	
	CR	SR	CR	SR
FINHEAR	30.81	1.95	25.96	1.68
ALSTM	16.77	0.75	8.97	0.41
Informer	8.54	1.22	5.06	0.82
PPO	15.46	1.67	10.33	1.01
FinCon	-8.17	-0.53	-13.32	-1.62
FinMem	19.24	1.43	15.09	1.19
CoT	22.07	1.32	16.14	1.28

Table 10: Impact of transaction costs on trading performance for the AAPL task. CR = Cumulative Return, SR = Sharpe Ratio. FINHEAR retains a significant advantage even under realistic cost assumptions.


A.4.7 Ablation

as shown in Table 11, based on the analysis of asset performance across the different time windows presented in the table, we observe a consistent pattern: various financial asset classes tend to achieve their peak relative performance over distinct time horizons. This differentiation is not coincidental but is intrinsically linked to their inherent asset characteristics and the temporal scale at which their core value drivers exert influence.


Assets whose value is highly contingent upon short-term market sentiment, immediate events, or rapid catalysts—such as gold (XAUUSD), which reacts swiftly to short-term uncertainty as a safe haven, or high-growth, high-volatility equities like Tesla (TSLA), prone to significant surges driven by recent favorable news or speculative momentum—tend to exhibit their most robust performance




(a) XAUUSD



(b) AAPL




(c) TSLA




(d) XOM


Figure 4: Cumulative trading performance of FinHEAR and other baselines on four individual assets (XAUUSD, AAPL, TSLA, XOM) from January 2023 to December 2023.




(a) XAUUSD



(b) AAPL



(c) TSLA



(d) XOM

Figure 5: Cumulative trading performance of FinHEAR and its ablation variants on four individual assets (XAUUSD, AAPL, TSLA, XOM) from January 2023 to December 2023

concentrated within shorter time windows. The performance of these assets is often characterized by high volatility in the short term, and sustaining such peak performance when averaged over extended time horizons can be challenging.

In contrast, optimal performance for some assets may manifest within a medium-term window, as exemplified by Apple (AAPL). The value drivers for this category of assets are typically more closely tied to factors such as product life cycles, consistent earnings growth, and medium-term industry trends. The impact of these factors requires a certain duration (typically several quarters to a few years) to fully materialize and be reflected in asset valuation. While extremely short windows may not adequately capture these cumulative effects, overly long horizons might introduce challenges such as growth deceleration or shifts in market cycles.

Furthermore, the performance of other assets is more likely to reach its optimum or show continuous improvement over longer time windows, with ExxonMobil (XOM) serving as an illustrative example. This class of assets is often closely linked to the long-term cycles of the macroeconomy, persistent trends in major commodity prices, or structural shifts in supply and demand. The evolution of these influencing factors is typically gradual and sustained, requiring several years to fully unfold. Over longer cycles, the impact of short-term market noise becomes diluted, and the asset’s underlying cyclical or fundamental long-term trend becomes dominant, enabling it to achieve superior cumulative performance.

From a financial theory perspective, this phenomenon highlights the concept of ‘Temporal Specificity of Asset Performance’ – where the dominant time scale of an asset’s core value drivers aligns effectively with the investment or observation horizon. Understanding this alignment is crucial for investors seeking to grasp the risk-return characteristics of various assets and to consider the temporal matching when constructing investment portfolios.

Table 11: Asset performance (%) over different time windows

Asset	Window = 3	Window = 5	Window = 7
XAUUSD	48.93	32.45	22.46
AAPL	29.15	30.81	23.56
TSLA	63.82	58.47	40.13
XOM	35.19	37.16	45.53

as shown in Figure 5a,Figure 5b,Figure 5cand Figure 5d, we present the cumulative trading performance of FinHEAR and its ablation variants on four individual assets: XAUUSD, AAPL, TSLA, and XOM. These ablation studies are conducted to evaluate the contribution of different key components of the FinHEAR model to its overall trading performance in single-asset markets. The figures illustrate the cumulative returns achieved by the full FinHEAR model compared to its variants with specific modules removed over the period from January 2023 to December 2023 for each respective asset, highlighting the effectiveness of each component.

A.4.8 Agent Modular Design and Prompt Engineering

This appendix provides a detailed exposition of representative prompt templates employed for the specialized agents within FinHEAR, serving to illustrate our approach to encoding distinct functionalities and expert reasoning paradigms. Each template is meticulously crafted to guide the Large Language Models (LLMs) in performing specific analytical and reasoning tasks, ensuring both fidelity to the intended behavioral profiles and structured output generation.

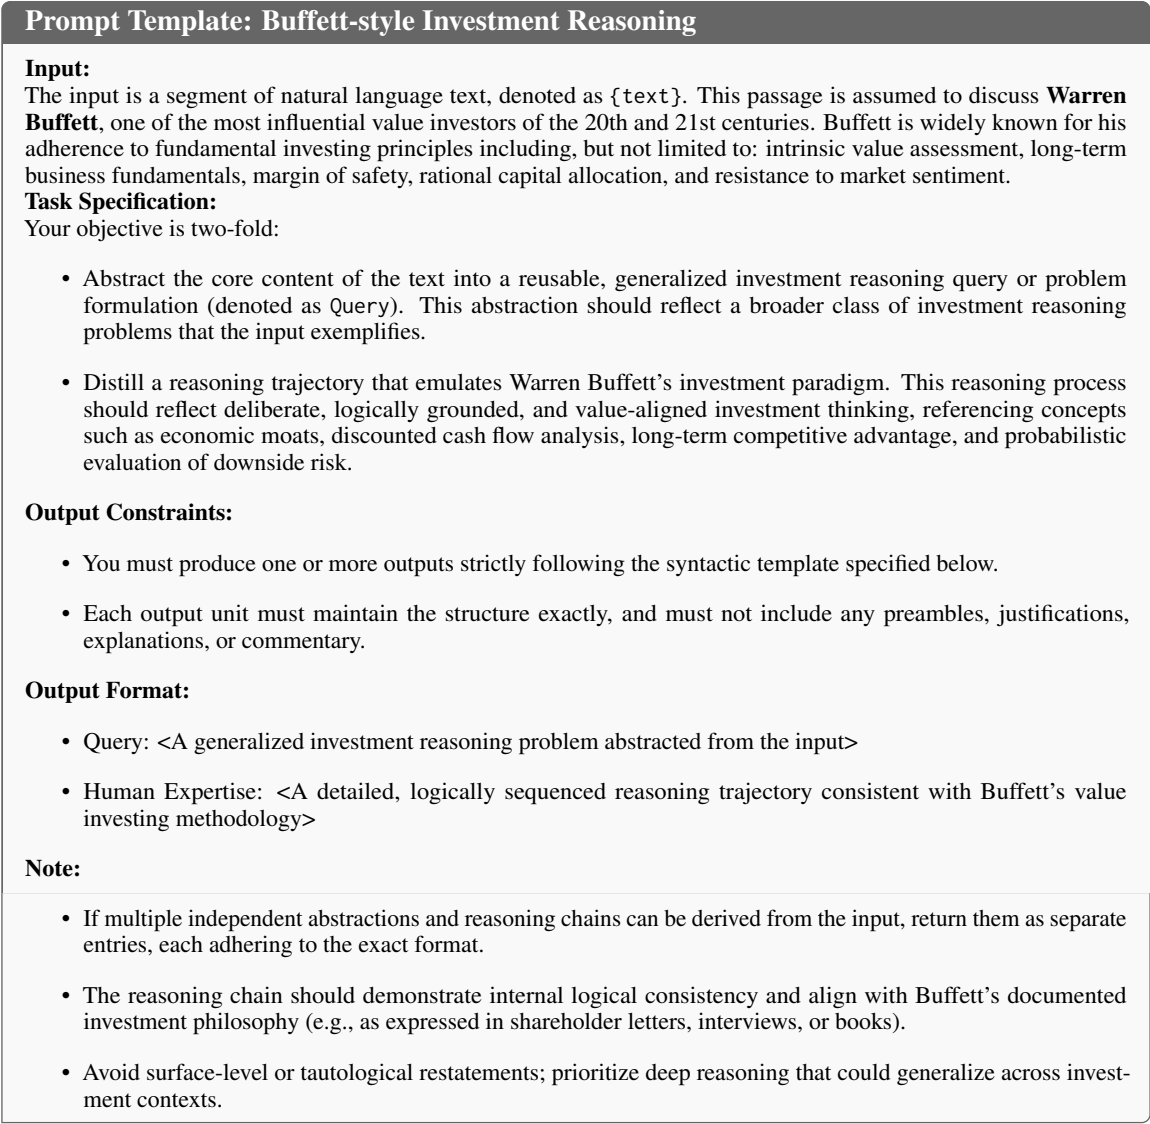


Figure 6: Prompt template for extracting generalized investment reasoning in the style of Warren Buffett.

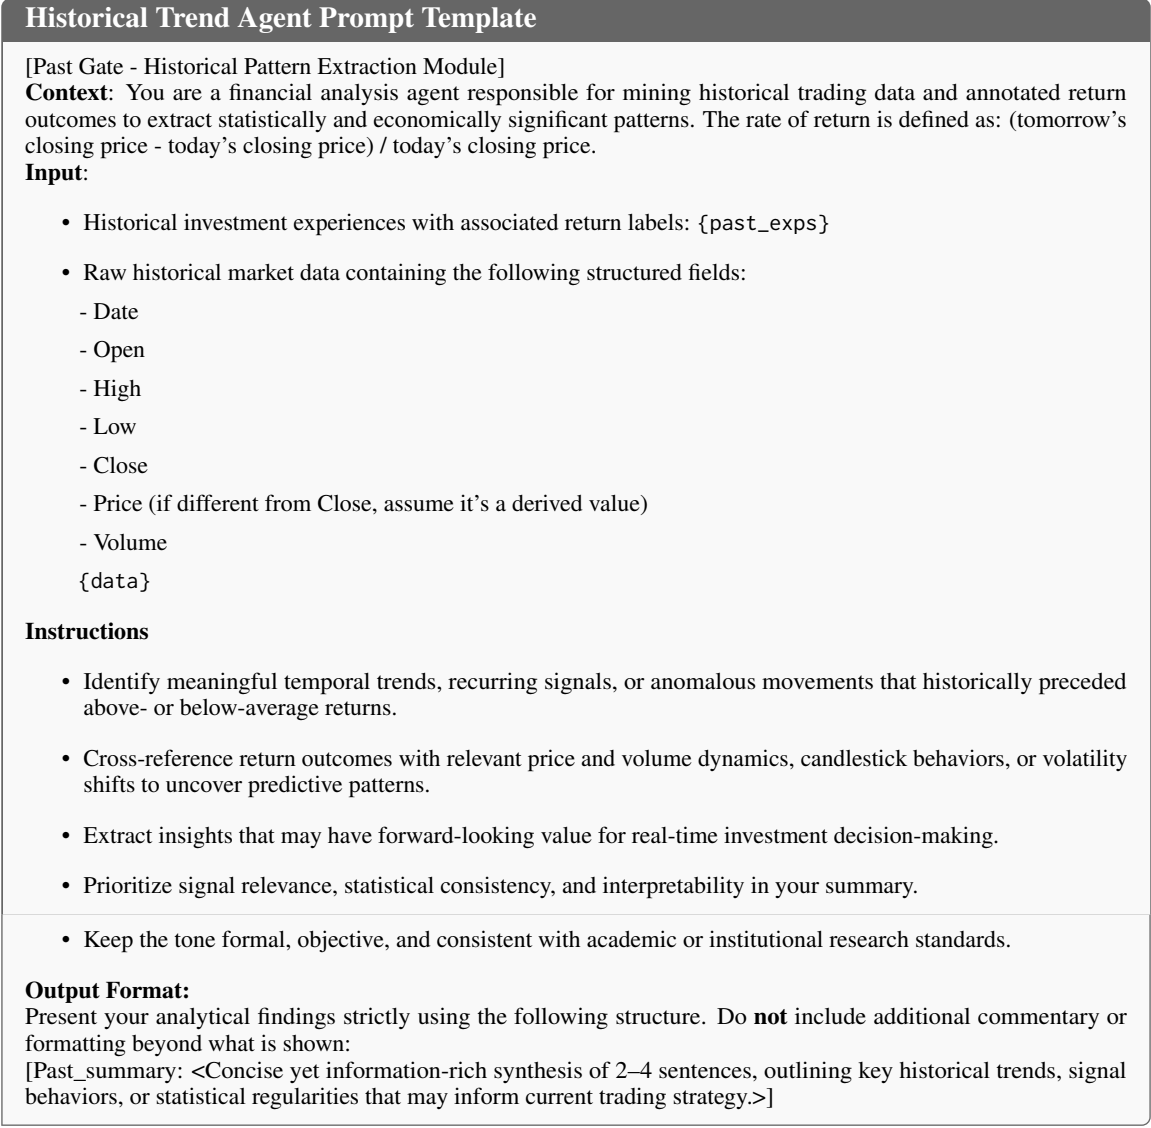


Figure 7: Prompt Template for Current Event Agent Financial Analysis

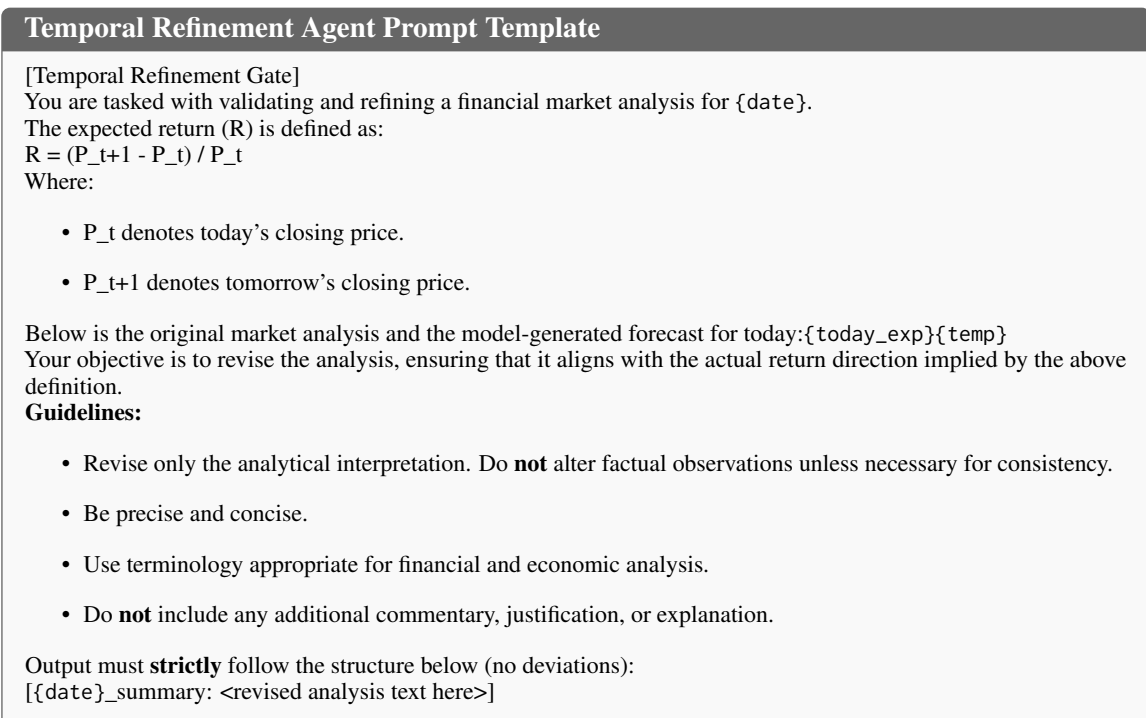


Figure 8: Prompt Template for Temporal Refinement Agent Financial Analysis and Decision

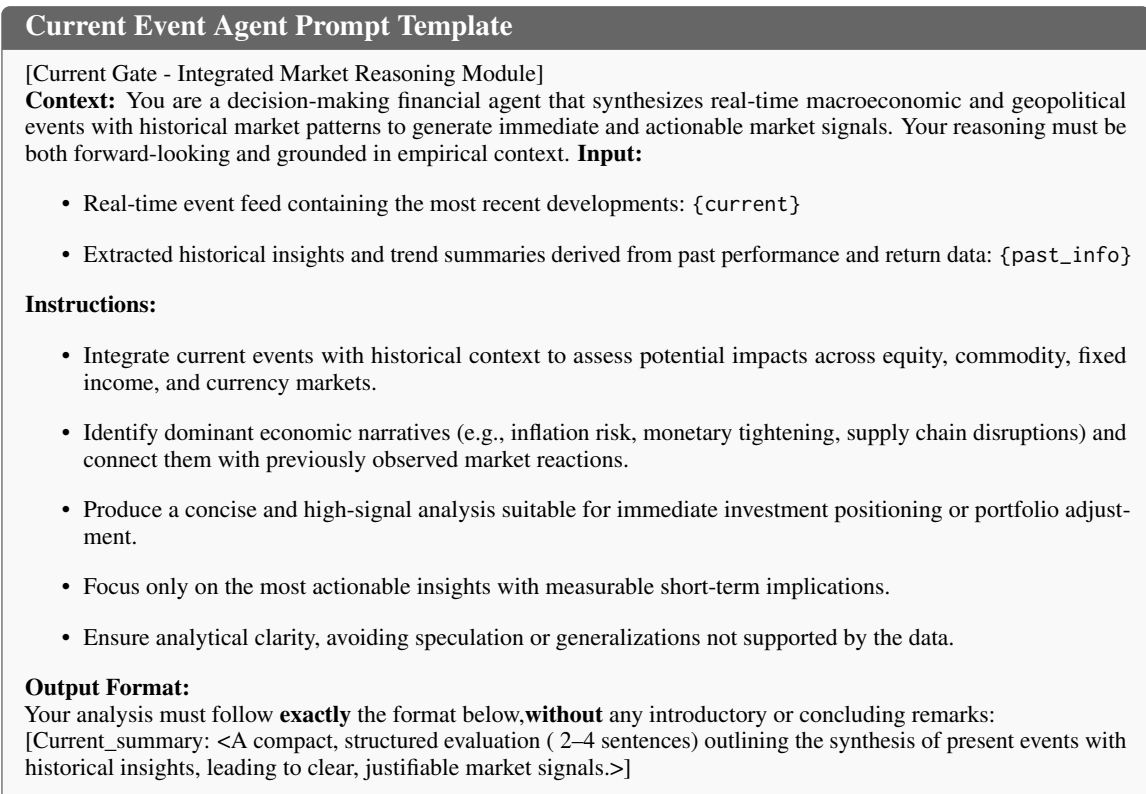


Figure 9: Prompt Template for Current Event Agent Financial Analysis

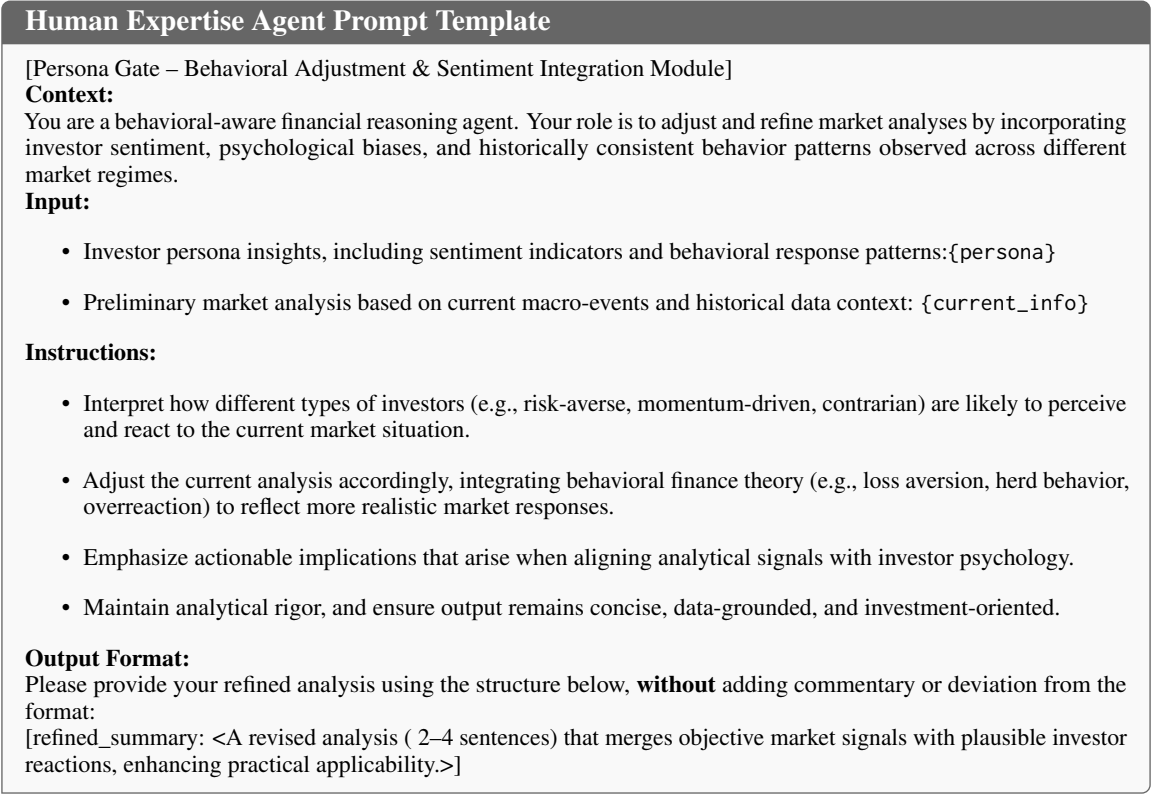


Figure 10: Prompt Template for Human Expertise Agent Financial Analysis

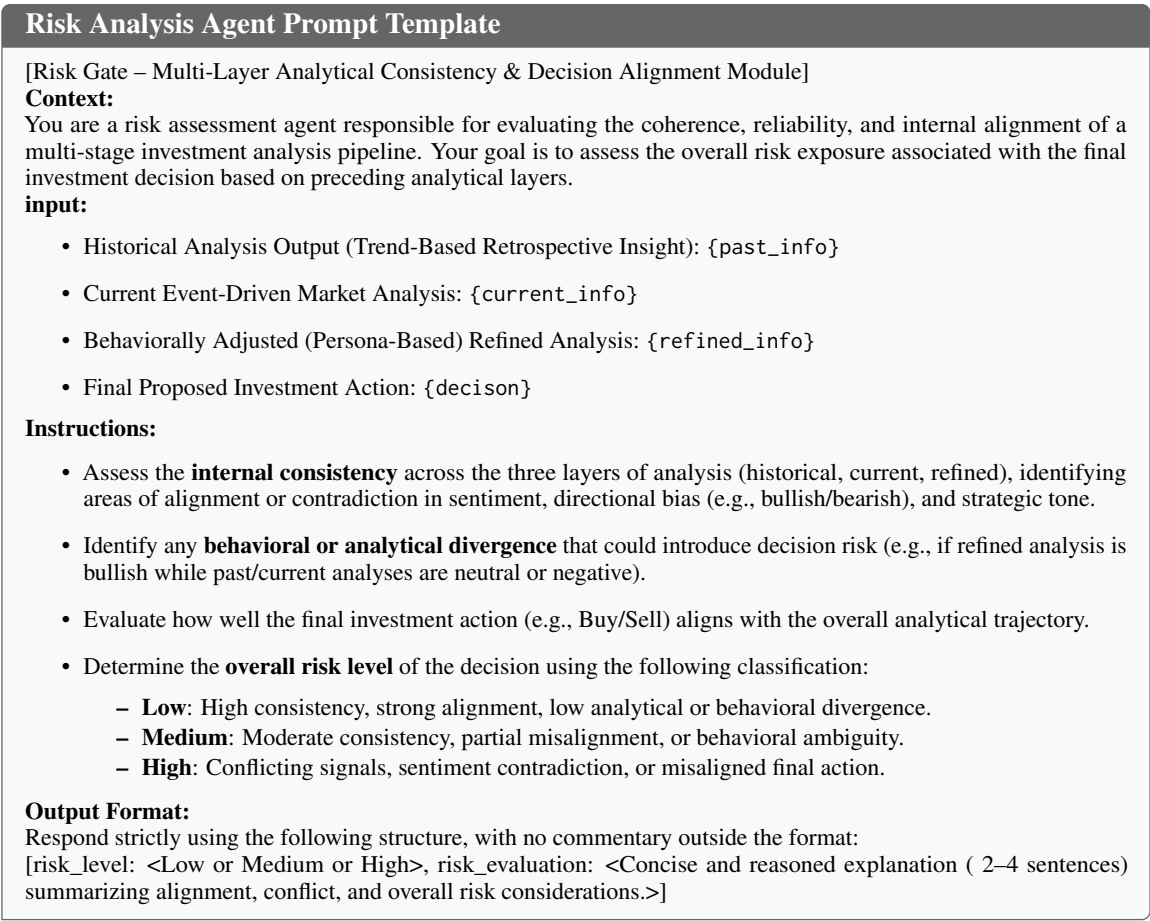


Figure 11: Prompt Template for Rise Analysis Agent Financial Analysis

Expert Alignment Agent Prompt Template

[Expert Alignment Gate – Multi-Agent Decision Validation Layer]

Role:

You are a centralized decision-making controller operating within a multi-agent investment advisory system. Each expert agent is specialized in a particular investment domain and asset class. Your responsibility is to assess the coherence, reliability, and macro-consistency of their respective recommendations in order to determine whether to adopt or reject their guidance.

Expert Profiles:

- **Warren Buffett:** A value investor with a long-term horizon, primarily focused on large-cap equities with durable competitive advantages. Assigned Asset: **Apple Inc. (AAPL)**
- **George Soros:** A macroeconomic and reflexivity-oriented investor. Specializes in trading global dislocations and regime shifts. Assigned Asset: **Gold Spot Market(XAUUSD)**
- **Peter Lynch:** A growth-oriented investor known for bottom-up stock selection and identifying fast-growing companies. Assigned Asset: **Tesla Inc. (TSLA)**
- **Benjamin Graham:** A fundamental value investor, originator of intrinsic valuation and margin-of-safety principles. Assigned Asset: **Exxon Mobil Corporation (XOM)**

Inputs Provided:

- **Macroeconomic Context** – A structured summary of the prevailing macro environment, including economic indicators, policy trends, market sentiment, and geopolitical signals: {macro_events}
- **Expert Recommendations** – Buy/Sell/Hold recommendations made by the four experts for their respective assets: {decisions}

Task Instructions:

- For each expert, analyze whether their recommendation is **contextually valid**, based on:
 - Their specific investment philosophy and historical behavioral traits
 - The macroeconomic landscape and potential asset sensitivity
 - Cross-asset interdependencies and systemic market risks
- Use the following decision rules:
 - Respond ****Yes**** if the expert's recommendation is well-aligned with both the macroeconomic context and their own philosophical model of decision-making.
 - Respond ****No**** if there is misalignment, excessive risk exposure, or evident inconsistency with current conditions.
- Your final output should reflect **discrete endorsement decisions**, not nuanced analysis. These decisions are intended to drive automated position sizing and portfolio response.

Output Format (strictly enforced):

Return a single-line output with no additional explanation, using the format below:

[Buffett:<Yes or No>, Soros:<Yes or No>, Lynch:<Yes or No>, Graham:<Yes or No>]

Figure 12: Prompt Template for Expert Alignment Agent Financial Analysis and Decision