Don't Forget the Last Trade: Integrating Episodic Recall into Memory-Driven Agents for High-Frequency Trading

Anonymous ACL submission

Abstract

Large language models (LLMs) have demonstrated strong capabilities in financial analysis and reasoning, inspiring recent advances in agent-based trading frameworks. While these systems show promise, prior approaches either emphasize long-horizon forecasting or operate as stateless analyzers, limiting their applicability to the demands of high-frequency trading (HFT). To address these gaps, we introduce META (Memory Enhanced Trading Agent), the first memory-augmented multi-agent framework for HFT. META integrates a family of specialized indicator agents (e.g., Trend, MACD, Stochastic, RSI, SMA, AVWAP, Heikin-Ashi) with a Decision Agent that fuses their reports, and a Memory Agent that retrieves and updates past trading episodes encoded as market state embeddings with outcomes and reflections. By recalling relevant experiences and adaptively reweighting signals under similar market regimes, META achieves improved directional accuracy and robustness under short-horizon evaluation. Our results demonstrate that episodic memory provides a powerful mechanism for regime-aware, interpretable, and low-latency decision-making in high-frequency trading.

1 Introduction

Building autonomous agents capable of reliable decision-making in complex, non-stationary environments remains a central challenge in machine learning (Weng, 2023; Sumers et al., 2023). High-frequency trading (HFT) exemplifies this difficulty: agents must process rapidly evolving, high-dimensional data under strict latency constraints, where profitable patterns can invert within minutes (Ding et al., 2024; Yang et al., 2023; Zhang et al., 2023a). Traditional quantitative systems rely on handcrafted logic for low-latency execution but collapse under regime shifts. Conversely, Large Language Model (LLM) agents demonstrate advanced reasoning ability (Wang et al., 2024b; Li et al., 2023b), yet remain *stateless*—each decision made in isolation, without learning from prior outcomes. This absence of experiential continuity limits adaptability in markets defined by recurrence and volatility.

We argue that progress in this domain requires agents endowed with **experiential memory**, capable of recognizing and reusing structural analogies between past and present market states (Yu et al., 2023; Ji et al., 2023). While dynamic memory frameworks such as Mem0 and MemGPT (Chhikara et al., 2025; Packer et al., 2023; Xu et al., 2025) enable long-term reasoning in text-based tasks, they fail to generalize to the stochastic, high-dimensional structure of financial time series. Unlike language, market data evolve as continuous numerical manifolds, demanding memory systems that operate over both symbolic abstractions and quantitative dynamics. This motivates a new direction: re-engineering episodic memory for real-time environments where state representations are not linguistic but numerical and non-stationary.

To this end, we introduce **META** (Memory Enhanced Trading Agent), a multi-agent cognitive architecture that reframes trading as a "Perceive–Recall–Synthesize" process. In the **Perception** stage, specialized Signal Agents extract interpretable technical features from market streams. In the **Recall** stage, a Memory Operator performs case-based retrieval from prior trades to identify historically analogous contexts. Finally, the **Synthesis** stage fuses real-time analysis with retrieved precedents to produce context-aware trading decisions. This process is powered by a **dual-stream state representation** that separates symbolic indicator vectors (for reasoning and retrieval) from normalized price–volume embeddings (for similarity search), balancing interpretability and numerical precision.

META is the first cognitive trading framework to embed retrieval-augmented episodic memory di-

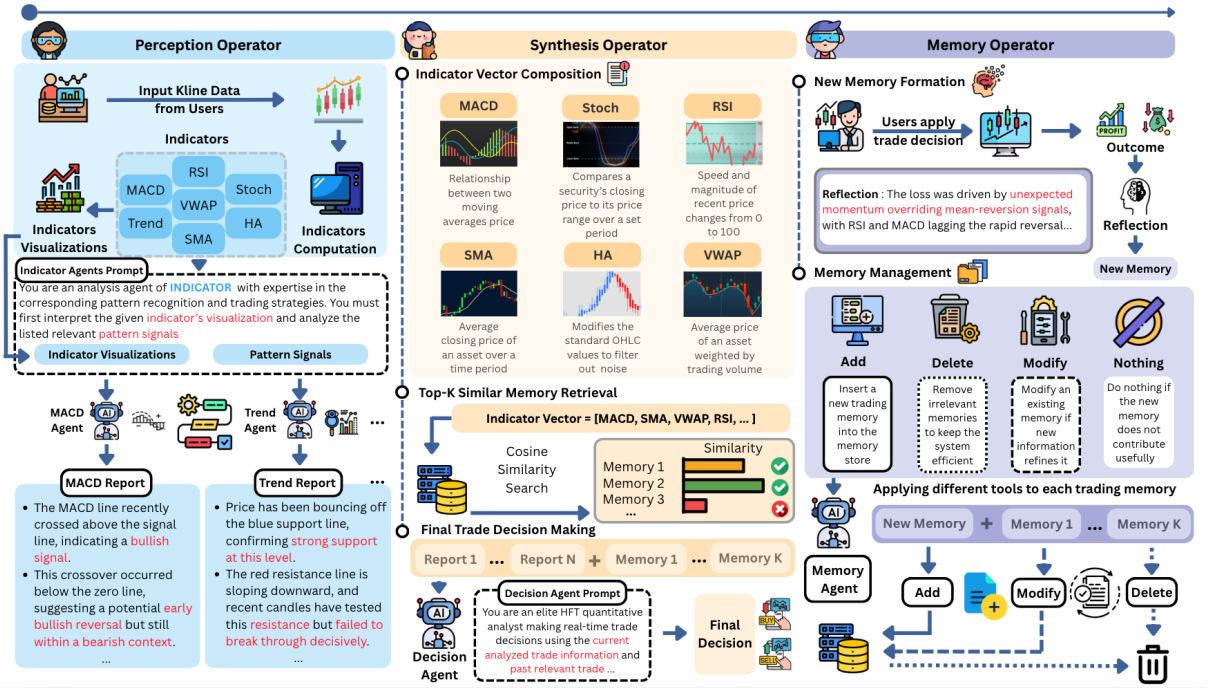

Figure 1: The META cognitive workflow, illustrating the interplay of its three core operators. The process unfolds in stages: (1) **Perception** (Φ): A distributed ensemble of Signal Agents encodes the market state in parallel. (2) **Synthesis** (Π): Real-time reports are fused with retrieved historical precedents to formulate a decision. (3) **Memory** (Ψ): The memory buffer is updated via a reflection-augmented consolidation cycle.

rectly into the high-frequency decision loop. Its memory operator extends dynamic memory principles (Chhikara et al., 2025; Li et al., 2023a; An et al., 2025) to structured numerical domains, enabling adaptive reasoning grounded in market feedback. Empirically, META achieves substantial gains in profitability and stability over stateless baselines, showing that coupling cognition with memory yields agents that not only react to markets—but *learn* from them.

2 Related Work

LLMs for Financial Decision-Making. Large Language Models (LLMs) are increasingly deployed as autonomous agents for financial trading, analysis, and alpha discovery (Ding et al., 2024; Yang et al., 2023; Zhang et al., 2023a; Wu et al., 2023). LLM-based trading agents (Weng, 2023; Sumers et al., 2023) generally follow three paradigms: leveraging sentiment from news and social media (Lopez-Lira and Tang, 2023; Kirtac and Germano, 2024; Zhang et al., 2024a), using reinforcement learning (RL) to refine policies through simulated trading (Ding et al., 2023; Koa et al., 2024), or employing explicit reasoning frameworks (Ji et al., 2023; Li et al., 2023b; Yu et al., 2023; Zhang et al., 2024b). Our work, META, advances the reasoning-driven paradigm with a specialized memory framework for technical analysis. It uses an "Indicator Vector" (e.g., MACD, RSI) to retrieve analogous historical scenarios via cosine similarity. A key contribution is our *Memory Update and Reflection* loop, which refines memory based on trade outcomes, allowing the agent to learn from the performance of technical patterns. Unlike systems that generate new alpha factors (Xiong et al., 2025; Wang et al., 2023b), META focuses on adaptive decision-making with established indicators.

Memory Mechanisms in Agents. Long-term reasoning is a central challenge for LLM agents. Early methods that appended full interaction histories to prompts were inefficient and scaled poorly (Hatalis et al., 2023; Liu et al., 2024; An et al., 2025). This spurred the development of external memory frameworks, including layered architectures, dynamic knowledge networks, and cognitively inspired systems (Packer et al., 2023; Xu et al., 2025; Lee et al., 2024; Salama et al., 2025; Wang et al., 2023a). However, persistent challenges include memory modules that are decoupled from the agent's reasoning policy (Li et al., 2023a; Yoon et al., 2024; Chhikara et al., 2025) and RL-based

systems that still rely on full-history accumulation without selective retention (Jin et al., 2025; Zheng et al., 2025). META addresses these limitations in quantitative trading by tightly integrating a structured, indicator-based memory with a reflection-driven decision cycle. By using cosine similarity retrieval and continuously updating its knowledge based on trade performance, META enhances adaptivity in non-stationary market environments.

3 Method

We propose **META** (**M**emory-**E**nhanced **T**rading **A**gent), a cognitive framework that decomposes the trading process into three modular operators: Perception (Φ), Memory (Ψ), and Synthesis (Π). Unlike conventional end-to-end black-box models, META introduces an explicit and auditable reasoning workflow. Its compositional design enables functional modularity, leverages non-parametric memory for generalization, and supports interpretable, retrieval-grounded decision making.

3.1 Problem Formulation

We cast trading as the problem of learning a *retrieval-augmented policy* π, which maps a market state $z_t \in \mathcal{S}$ and a time-varying episodic memory $\mathcal{M}_t$ to a structured decision d_t. The objective is to learn an optimal policy $\pi^\star$ that maximizes the expected cumulative return:

$$J(\pi) = \mathbb{E}\left[\sum_{t=0}^{T} \gamma^t \, o_t(d_t)\right], \pi^\star = \arg\max_{\pi} J(\pi),$$

where o_t denotes the outcome from decision d_t, and $\gamma \in (0, 1]$ is a temporal discount factor.

At the heart of META lies the hypothesis that optimal trading policies are *functionally decomposable*. We define π as the composition of two interpretable modules:

$$d_t = \Pi\big(\underbrace{\Phi(z_t)}_{\text{Perception}}, \underbrace{\Psi(z_t, \mathcal{M}_t)}_{\text{Memory}}\big).$$

This structured decomposition confers three main benefits: (i) **specialized processing** of heterogeneous financial inputs; (ii) **retrieval-augmented generalization** through episodic memory; and (iii) **transparent inference chains** that elucidate the agent's decision logic (Fig. 1).

3.2 Bifurcated State Representation

Financial markets exhibit both high-level structural regimes and fine-grained temporal fluctuations. To model this duality, META employs a *bifurcated state representation* that jointly captures symbolic abstractions and numerical microstructure. From a sliding window of observations $X_t = \{(O, H, L, C, V)_{t-i}\}_{i=0}^{W-1}$, we construct:

$$z_t = \left[z_t^{(\text{sym})} \,\|\, z_t^{(\text{num})}\right],$$

where $\|$ denotes concatenation.

The **symbolic component** $z_t^{(\text{sym})} \in \mathbb{R}^M$ encodes several interpretable, domain-specific features—such as momentum, volatility, and mean-reversion indicators—using compact heuristic transformations. This channel supports logical and rule-based reasoning, providing the cognitive substrate for high-level pattern recognition and symbolic inference.

The **numerical component** $z_t^{(\text{num})} \in \mathbb{R}^{5W}$ preserves the normalized, high-fidelity raw market microstructure (open, high, low, close, volume). It enables precise retrieval through vector similarity, capturing minute market configurations often lost in purely symbolic abstractions.

This bifurcated encoding allows the architecture to reason at two complementary levels: the symbolic stream offers interpretability and abstraction, while the numerical stream ensures precision and adaptability. By leveraging both representations, META bridges symbolic and subsymbolic reasoning, a capability central to hybrid systems.

3.3 Perception Operator (Φ)

The perception operator Φ implements a distributed ensemble of *Signal Agents* $\{A^m\}_{m=1}^{M}$, where each agent embodies a domain heuristic or learned detector (see Appendix B for the detailed mathematical formulation of each indicator). Each A^m processes the market state z_t to produce a compact report:

$$r_t^m = A^m(z_t).$$

The collective output forms a perceptual set $R_t = \{r_t^m\}_{m=1}^{M}$, in which each report r_t^m encapsulates both quantitative metrics (e.g., indicator strength) and qualitative semantics (e.g., *"bullish divergence," "breakout detected"*).

This design provides two crucial benefits. First, it modularizes perception—each agent can be refined or replaced independently, supporting continuous integration of new trading signals. Second, it yields structured, interpretable percepts, allowing higher modules to perform symbolic reasoning without operating directly on raw data. This structure mirrors biological cognition, where perceptual abstractions precede higher-level inference.

3.4 Memory Operator (Ψ)

The memory operator Ψ implements *case-based analogical reasoning*. The episodic memory $\mathcal{M} = \{m_i\}_{i=1}^N$ stores discrete trading experiences, where each entry $m_i = (z_i, a_i, o_i, \rho_i)$ logs the state, action, outcome, and an optional textual reflection summarizing the context and lesson learned.

Retrieval. Given the current state z_t, Ψ identifies historically similar contexts by computing cosine similarity over the numerical embeddings:

$$\mathrm{sim}(z_t, z_i) = \frac{(z_t^{(\mathrm{num})})^\top z_i^{(\mathrm{num})}}{\|z_t^{(\mathrm{num})}\| \|z_i^{(\mathrm{num})}\|}.$$

A top-k selection of analogous cases forms the retrieved neighborhood:

$$\mathcal{N}_t = \{(m_i, s_i) \mid s_i = \mathrm{sim}(z_t, z_i) \geq \tau\}_{\text{Top-}k}.$$

These retrieved cases constitute the agent's "episodic evidence," grounding its current reasoning in empirical precedent.

Consolidation. To preserve computational efficiency and diversity, Ψ employs a dynamic curation operator $\mathrm{Op}(m_t, \mathcal{M})$ that manages the lifecycle of memories. Low-utility or redundant entries are pruned, while novel, high-impact experiences are retained. This continual abstraction transforms $\mathcal{M}$ into a structured, evolving repository of archetypal scenarios—a non-parametric analog of synaptic consolidation in biological systems.

3.5 Synthesis Operator (Π)

The synthesis operator Π employs a LLM to integrate perceptual summaries R_t and retrieved episodes $\mathcal{N}_t$ into a structured decision output:

$$d_t \triangleq (a_t, \kappa_t, \mathcal{E}_t, \rho_t),$$

where a_t denotes the action (e.g., `LONG`), κ_t is the confidence score, $\mathcal{E}_t \subseteq R_t \cup \mathcal{N}_t$ is the subset of selected evidence, and ρ_t is a concise natural language rationale justifying a_t.

We enforce an **evidence-first, constrained decoding** process: the LLM must first select $\mathcal{E}_t$ from the available context, and only then generate the rationale ρ_t. This separation enforces grounded reasoning and produces an auditable explanation chain from input evidence to trading action.

4 Experiments and Results

4.1 Experimental Setup

To evaluate the multi-agent framework of META, we perform experiments on a asset consisting of five representative financial assets constructed by prior work Xiong et al. (2025): Crude Oil Futures (CL), S&P 500 E-mini Futures (ES), Nasdaq-100 Futures (NQ), Invesco QQQ ETF (QQQ), and Bitcoin (BTC). Each dataset consists of 4-hour OHLCV (Open, High, Low, Close, Volume) sequences sampled over multi-year horizons, capturing diverse market regimes and volatility patterns. Following the experimental protocol of Xiong et al. (2025), each asset is partitioned into non-overlapping windows of 100 candlesticks, with the final three windows reserved for validation.

For the memory configuration, META employs the top-k and similarity-threshold retrieval strategy. Unless otherwise stated, k is set to 5 and the similarity threshold τ to 0.7. This configuration balances computational efficiency and contextual relevance, enabling low-latency yet adaptive operation suited for high-frequency trading environments.

As baselines, we include the QUANTAGENT framework (Xiong et al., 2025)—a multi-agent trading system without long-term memory—and a random decision generator that selects between LONG and SHORT uniformly at random while sharing the same evaluation segments and reporting format. All methods are evaluated under consistent metrics and experimental conditions.

Evaluation. We adopt the evaluation protocol proposed by Xiong et al. (2025) to ensure comparability. The primary metric is **directional accuracy** (α), which quantifies the correctness of LONG/SHORT forecasts relative to subsequent price movements. In addition, we report four complementary return-based metrics from Xiong et al. (2025) that jointly assess profitability, robustness, and memory relevance: Average cumulative return correlation (R_{cc}) Similarity-weighted return alignment (R_{sim}) Maximum realized gain ($R_{\max}$) Maximum drawdown ($R_{\min}$) Detailed definitions, computation procedures, and theoretical motivation for these metrics are provided in Appendix D.

4.2 Main Results

Table 1 presents a comprehensive comparison between META, QUANTAGENT, and a random trading policy across five assets. The results demon-

Method	Accuracy α (% ↑)	$\Delta\alpha$ (% ↑)	R_{cc} ↑	$R_{\max}$ ↑	$R_{\min}$ ↑	R_{sim} ↑
CL						
+ Random	49.0	–	-0.312	0.999	-1.384	-0.255
+ QuantAgent	**57.7**	**+17.8%**	**-0.195**	**1.181**	-1.202	**-0.133**
+ META (Ours)	56.0	+6.1%	-0.484	1.024	**-1.187**	-0.433
ES						
+ Random	41.3	–	0.006	0.560	-0.539	0.006
+ QuantAgent	55.0	+33.2%	0.179	**0.613**	-0.485	0.179
+ META (Ours)	**64.0**	**+55.0%**	**0.278**	0.584	**-0.431**	**0.250**
NQ						
+ Random	41.3	–	-0.033	0.717	-0.736	-0.033
+ QuantAgent	53.3	+29.1%	0.078	**0.747**	-0.705	**0.078**
+ META (Ours)	**62.0**	**+50.1%**	**0.089**	0.631	**-0.667**	-0.029
QQQ						
+ Random	39.7	–	-0.265	0.966	-1.038	-0.252
+ QuantAgent	59.7	+50.4%	0.189	1.052	**-0.952**	**0.193**
+ META (Ours)	**61.5**	**+54.9%**	**0.627**	**1.094**	-1.102	-0.029
BTC						
+ Random	44.3	–	-0.259	1.115	-1.366	-0.277
+ QuantAgent	50.7	+14.5%	0.081	1.232	-1.249	**0.004**
+ META (Ours)	**56.0**	**+26.4%**	**0.120**	**1.296**	**-0.980**	-0.040

Table 1: Performance comparison across assets. Accuracy (α), relative improvement ($\Delta\alpha$), correlation R_{cc}, maximum and minimum returns ($R_{\max}$, $R_{\min}$), and similarity metric (R_{sim}). Red indicates better performance, Green indicates worse, and black numbers denote baseline or neutral results.

strate that META consistently outperforms both baselines in directional accuracy and return-based metrics, confirming the effectiveness of its memory-augmented decision process.

In terms of accuracy (α), META achieves the highest performance on all assets except Crude Oil (CL), where it performs competitively with QUANTAGENT. Notably, on S&P 500 E-mini (ES), accuracy rises from 55.0% in QUANTAGENT to 64.0% with META, corresponding to a relative improvement of +55.0%. Similar gains are observed on Nasdaq-100 (NQ) and Bitcoin (BTC), where accuracy improves by +50.1% and +26.4%, respectively. Even for more stable equity markets such as QQQ, META maintains a marginal but consistent edge over QUANTAGENT (61.5% vs. 59.7%). Beyond directional accuracy, the return-based metrics further corroborate the effectiveness of META. The model consistently achieves the strongest cumulative return correlations (R_{cc}) on ES (0.278), QQQ (0.627), and BTC (0.120), indicating that its predicted trade directions align more reliably with realized price movements compared to prior baselines. Moreover, META demonstrates notably improved downside resilience, achieving the least negative drawdowns ($R_{\min}$) across all assets except QQQ, where it remains competitive. This robustness reflects the stabilizing effect of the memory-augmented reasoning mechanism, which allows the agent to recall profitable patterns while avoiding previously adverse conditions.

4.3 Text Embeddings vs. Indicator Vectors

We compare META's performance under two retrieval modalities: text embeddings derived from agent reports versus indicator vectors composed of technical features and normalized OHLCV data.As shown in Figure 3, indicator-vector retrieval consistently outperforms text-embedding retrieval. On ES futures, accuracy improves from 56.0% to 64.0%, and on QQQ from 50.0% to 60.0%. Even for more stable markets such as CL and BTC, indicator retrieval yields gains of 2–6 percentage points. Overall, indicator-based retrieval achieves 2–10% higher accuracy than text embeddings.

Beyond accuracy, the two modalities exhibit distinct retrieval geometries. Text embeddings produce uniformly high similarity scores (0.90–0.91) with minimal variance ($\approx$ 0.0002–0.0003), implying that semantically similar reports cluster together regardless of market context—limiting discriminative power. In contrast, indicator vectors yield lower mean similarities (0.64–0.70) but higher variance ($\approx$ 0.0026–0.0033), indicating

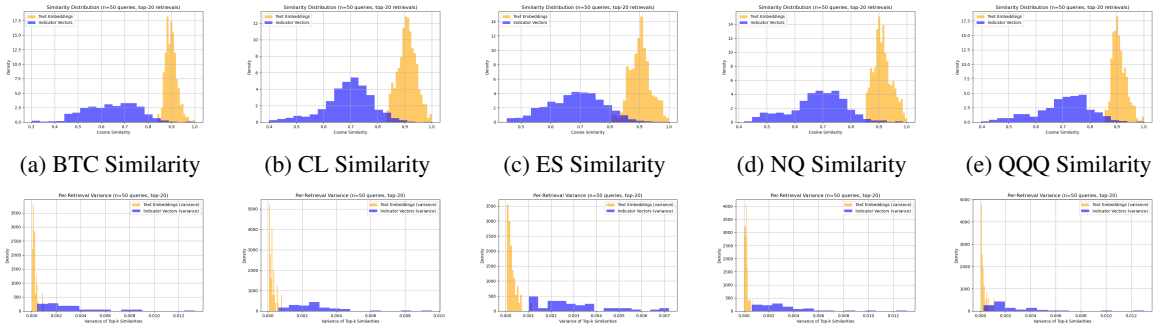

(a) BTC Similarity (b) CL Similarity (c) ES Similarity (d) NQ Similarity (e) QQQ Similarity

(f) BTC Variance (g) CL Variance (h) ES Variance (i) NQ Variance (j) QQQ Variance

Figure 2: Similarity and variance distributions for text embeddings vs. indicator vectors across all assets. Each asset is shown with two plots: top for similarity distribution and bottom for variance distribution.

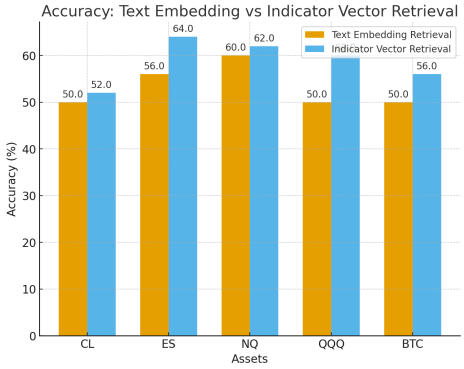

Figure 3: Accuracy comparison of text-embedding vs. indicator-vector retrieval.

Asset	Similarity (Mean)		Variance (Mean)	
	Text	Indic.	Text	Indic.
BTC	0.9002	0.6428	0.0002	0.0031
CL	0.9089	0.6905	0.0002	0.0028
ES	0.9036	0.6853	0.0003	0.0030
NQ	0.9139	0.6826	0.0002	0.0033
QQQ	0.9095	0.7011	0.0003	0.0026

Table 2: Comparison of mean similarity and variance between text embeddings and indicator vectors. Indicator vectors show lower similarity but higher variance, leading to more discriminative retrieval.

sharper separation between relevant and irrelevant precedents. This richer similarity spread provides a stronger ranking gradient, improving retrieval selectivity—crucial for high-frequency trading, where noisy matches can quickly degrade performance. These findings confirm META's design choice to ground memory retrieval in structured indicator vectors rather than textual embeddings.

4.4 Memory Retrieval Guidance

We analyze the role of the top-k and similarity threshold (τ) hyperparameters in controlling memory retrieval quality and efficiency across assets. These experiments validate META's design choice of combining both mechanisms for stable and adaptive recall.

Retrieved counts at fixed threshold. At a fixed threshold (e.g., $\tau = 0.7$), the number of retrieved memories per query (Figure 4, top) increases steadily as the memory base grows. Although individual counts fluctuate, the rolling average rises over time, showing that thresholding functions as a quality filter—restricting retrieval early when memory is sparse, and gradually broadening it as more relevant experiences accumulate.

Average top-k similarity at fixed k. With k fixed (e.g., $k = 10$), the average similarity of retrieved memories (Figure 4, middle) improves as the memory expands, stabilizing at a high level. This demonstrates that top-k selection maintains focus on the most relevant precedents while preventing uncontrolled growth in reference size.

Joint effect of k and τ. Heatmaps of average retrieval counts across k–τ combinations (Figure 4, bottom) reveal their complementary dynamics. At low τ, increasing k rapidly enlarges the retrieved set, while very high τ collapses retrieval to near zero. Balanced behavior occurs around moderate settings ($\tau = 0.6$–0.7, $k = 5$–15), yielding both sufficient coverage and high reliability.

4.5 Learning Dynamics with Memory Accumulation

We examine how META's performance evolves as its memory base grows. Figure 5 shows the moving-average directional accuracy. The analysis starts at 300 memories, since the initial 0–300 trades rely primarily on perception and reasoning modules

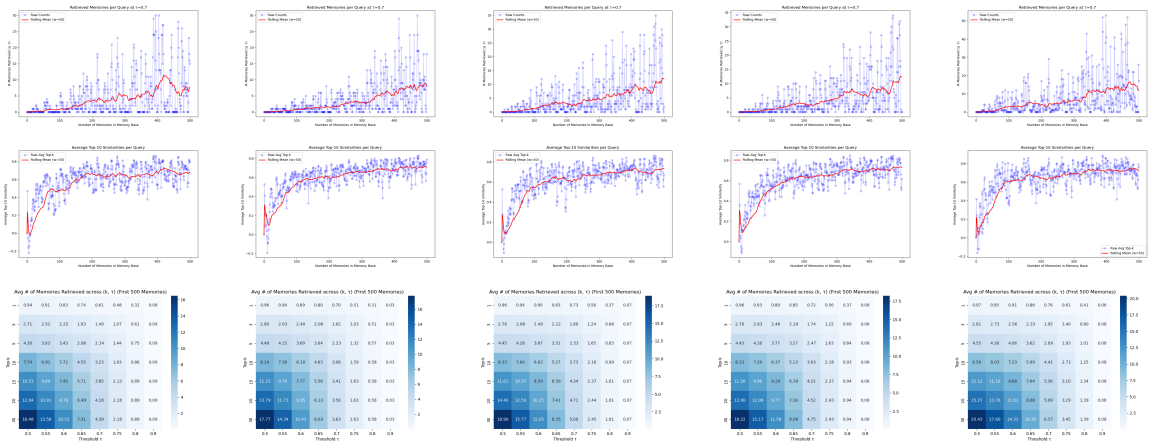

(a) BTC (b) CL (c) ES (d) NQ (e) QQQ

Figure 4: Top-k and threshold retrieval analysis. Each asset is shown with three plots: (top) # memories retrieved vs. memory base size at $\tau = 0.7$, (middle) average top-k similarities vs. memory base size at $k = 10$, and (bottom) heatmap of average retrieval counts over k and τ.

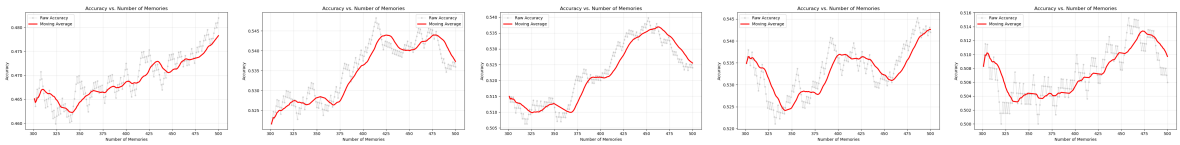

(a) CL (b) ES (c) NQ (d) QQQ (e) BTC

Figure 5: Accuracy versus number of memories. Each curve shows the moving average of directional accuracy as META's memory base expands. The analysis begins at 300 memories, corresponding to the point where episodic retrieval starts to meaningfully influence decision-making.

rather than episodic recall; early accuracy variations thus reflect exploratory behavior rather than learned adaptation. Beyond this phase, all assets exhibit a clear upward trend in accuracy, demonstrating consistent performance gains from memory accumulation. Between 300–350 trades, moderate volatility appears as retrieval stabilizes and the memory base diversifies. As memory density increases further, accuracy steadily improves and variance declines, confirming that broader episodic coverage enhances contextual reasoning and decision precision. Notably, assets with higher volatility such as BTC and NQ display sharper accuracy gains, highlighting the value of episodic retrieval in turbulent markets. In contrast, more stable assets like CL and ES show smoother yet consistent improvements, underscoring META's robustness across diverse market conditions.

4.6 Robustness Under Varying Market Conditions

We evaluate META's stability across four representative market regimes—*uptrend*, *downtrend*, *steady*, and *sudden change*—to assess its adaptability under dynamic trading environments. As shown in Figure 6, performance remains highly consistent across regimes, with accuracy variations confined to within 4–6% across all assets.

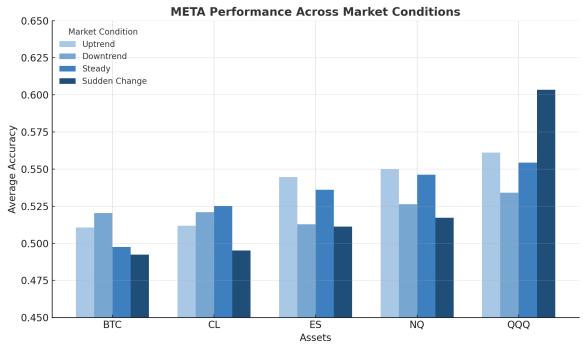

Figure 6: Performance of META under different market conditions across assets.

Across the five assets (BTC, CL, ES, NQ, and QQQ), META achieves its greatest stability in steady markets, maintaining an average accuracy above 0.53. This reflects the model's ability to exploit recurring structural patterns through memory retrieval. In both *uptrend* and *downtrend*

Asset	Method	Accuracy α (% ↑)	$\Delta\alpha$ (↑)
CL	w/o Mem	52.5	–
	w/ Mem	**56.0**	**+3.5**
ES	w/o Mem	50.0	–
	w/ Mem	**64.0**	**+14.0**
NQ	w/o Mem	40.0	–
	w/ Mem	**62.0**	**+22.0**
QQQ	w/o Mem	55.0	–
	w/ Mem	**61.5**	**+6.5**
BTC	w/o Mem	48.0	–
	w/ Mem	**56.0**	**+8.0**

Table 3: Ablation study of memory system.

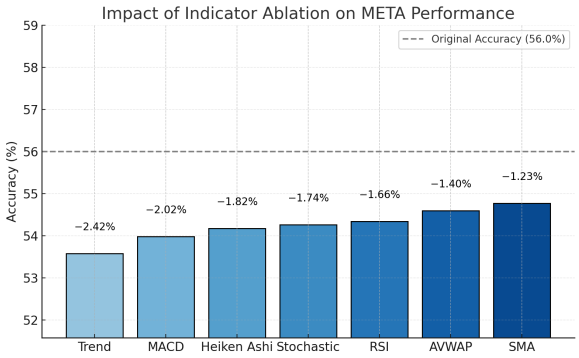

Figure 7: Impact of removing each technical indicator on META's directional accuracy. Larger decreases indicate greater importance of that indicator to overall system stability and predictive accuracy.

phases, comparable accuracy levels indicate that META generalizes directional dependencies effectively without overfitting to a particular bias. Importantly, the framework demonstrates strong resilience during sudden price shocks—most notably in QQQ (0.60 accuracy)—where conventional models typically degrade. This robustness stems from META's retrieval-guided decision process, which draws on structurally similar high-volatility episodes rather than short-term fluctuations, ensuring stability under abrupt regime shifts.

4.7 Memory System Ablation

We perform an ablation study to isolate the impact of the memory module by comparing META's performance with and without memory across five assets (Table 3). Across all settings, incorporating memory yields consistent and substantial gains in directional accuracy.

For high-volatility assets such as ES and NQ, the improvements are especially pronounced: accuracy increases from 50.0% to 64.0% (+14.0 points) on ES and from 40.0% to 62.0% (+22.0 points) on NQ. These results highlight the value of episodic retrieval in distinguishing subtle recurring patterns that short-term indicators alone fail to capture. Steadier assets also benefit, with CL improving from 52.5% to 56.0% (+3.5 points) and QQQ from 55.0% to 60.0% (+5.0 points). Even BTC, despite its noisy structure, shows an +8.0 point gain.

Overall, the consistent improvements across both volatile and stable markets confirm that the memory mechanism enhances predictive robustness and contextual reasoning, enabling META to adapt effectively to diverse trading conditions.

4.8 Indicator Contributions

To analyze the internal dynamics of the META framework, we perform an ablation study assessing the contribution of each technical indicator to overall trading performance. Each indicator—**MACD**, **AVWAP**, **Trend**, **RSI**, **Heiken Ashi**, **Stochastic**, and **SMA**—is removed individually, and the resulting change in directional accuracy quantifies its marginal impact on decision quality.

As shown in Figure 7, removing any single indicator reduces accuracy, confirming that all components contribute meaningfully to META's perception of market dynamics. **Trend** and **MACD** yield the largest declines (-2.42% and -2.02%), reflecting their central role in capturing mid- to long-term momentum. **Heiken Ashi** and **Stochastic** produce moderate drops (-1.82% and -1.74%), highlighting their complementary function in smoothing short-term volatility and detecting reversals. In contrast, **AVWAP** and **SMA** show smaller effects (-1.40% and -1.23%), suggesting partial redundancy with other trend-based indicators.

5 Conclusion

We present META, a memory-augmented trading framework that integrates multi-agent market perception with reflective reasoning. By combining specialized Indicator Agents, a decision fusion mechanism, and a dynamic Memory Operator, META enables adaptive and interpretable trading decisions in a zero-shot setting. Extensive experiments across diverse assets show that META consistently outperforms baselines, validating the effectiveness of memory-driven contextual recall for high-frequency decision making.

Limitations

While META demonstrates notable improvements in short-horizon trading accuracy and adaptability, several limitations remain. The current framework relies on handcrafted technical indicators for perceptual grounding, which may restrict scalability to more diverse market structures or alternative data modalities. Although the episodic memory improves contextual reasoning, its similarity-based retrieval remains static; future work could explore adaptive or learned retrieval metrics that evolve with market regimes. In addition, our evaluation focuses on retrospective simulation under historical data rather than live market execution, and thus omits slippage, transaction costs, and real-time latency constraints. Addressing these limitations would move META closer to a fully autonomous, production-ready cognitive trading agent.

References

An, C., Zhang, J., Zhong, M., Li, L., Gong, S., Luo, Y., Xu, J., and Kong, L. (2025). Why does the effective context length of LLMs fall short? In *Proceedings of the International Conference on Learning Representations (ICLR)*.

Chhikara, P., Khant, D., Aryan, S., Singh, T., and Yadav, D. (2025). Mem0: Building production-ready AI agents with scalable long-term memory. *arXiv preprint arXiv:2504.19413*.

Ding, H., Li, Y., Wang, J., and Chen, H. (2024). Large language model agent in financial trading: A survey.

Ding, Y., Jia, S., Ma, T., Mao, B., Zhou, X., Li, L., and Han, D. (2023). Integrating stock features and global information via large language models for enhanced stock return prediction.

Hatalis, K., Christou, D., Myers, J., Jones, S., Lambert, K., Amos-Binks, A., Dannenhauer, Z., and Dannenhauer, D. (2023). Memory matters: The need to improve long-term memory in LLM-agents. In *Proceedings of the AAAI Symposium Series*, 2:277–280.

Ji, Z., Yu, T., Xu, Y., Lee, N., Ishii, E., and Fung, P. (2023). Towards mitigating hallucination in large language models via self-reflection.

Jin, B., Zeng, H., Yue, Z., Yoon, J., Arik, S., Wang, D., Zamani, H., and Han, J. (2025). Search-R1: Training LLMs to reason and leverage search engines with reinforcement learning. *arXiv preprint arXiv:2503.09516*.

Kirtac, K. and Germano, G. (2024). Sentiment trading with large language models. *Finance Research Letters*, 62:105227.

Koa, K. J., Ma, Y., Ng, R., and Chua, T.-S. (2024). Learning to generate explainable stock predictions using self-reflective large language models.

Lee, K.-H., Chen, X., Furuta, H., Canny, J., and Fischer, I. (2024). A human-inspired reading agent with gist memory of very long contexts.

Li, Y., Dong, B., Guerin, F., and Lin, C. (2023a). Compressing context to enhance inference efficiency of large language models. In *EMNLP 2023*, pp. 6921–6935.

Li, Y., Yu, Y., Li, H., Chen, Z., and Khashanah, K. (2023b). TradingGPT: Multi-agent system with layered memory and distinct characters for enhanced financial trading performance.

Liu, N. F., Lin, K., Hewitt, J., Paranjape, A., Bevilacqua, M., Petroni, F., and Liang, P. (2024). Lost in the middle: How language models use long contexts. *Transactions of the ACL*, 12:157–173.

Lopez-Lira, A. and Tang, Y. (2023). Can ChatGPT forecast stock price movements? Return predictability and large language models.

Packer, C., Fang, V., Patil, S., Lin, K., Wooders, S., and Gonzalez, J. (2023). MemGPT: Towards LLMs as operating systems.

Salama, R., Cai, J., Yuan, M., Currey, A., Sunkara, M., Zhang, Y., and Benajiba, Y. (2025). MemInsight: Autonomous memory augmentation for LLM agents. *arXiv preprint arXiv:2503.21760*.

Sumers, T. R., Yao, S., Narasimhan, K., and Griffiths, T. L. (2023). Cognitive architectures for language agents.

Wang, B., Liang, X., Yang, J., Huang, H., Wu, S., Wu, P., Lu, L., Ma, Z., and Li, Z. (2023a). Enhancing large language models with self-controlled memory framework. *arXiv preprint arXiv:2304.13343*.

Wang, S., Yuan, H., Ni, L. M., and Guo, J. (2024b). QuantAgent: Seeking holy grail in trading by self-improving large language model. *arXiv preprint arXiv:2402.03755*.

Wang, S., Yuan, H., Zhou, L., Ni, L. M., Shum, H.-Y., and Guo, J. (2023b). AlphaGPT: Human–AI interactive alpha mining for quantitative investment. *arXiv preprint arXiv:2308.00016*.

Weng, L. (2023). LLM-powered autonomous agents.

Wu, D., Wang, H., Yu, W., Zhang, Y., Chang, K.-W., and Yu, D. (2025). LongMemEval: Benchmarking chat assistants on long-term interactive memory. In *International Conference on Learning Representations (ICLR)*.

Wu, S., Irsoy, O., Lu, S., Dabravolski, V., Dredze, M., Gehrmann, S., Kambadur, P., Rosenberg, D., and Mann, G. (2023). BloombergGPT: A large language model for finance.

Xie, Q., Han, W., Zhang, X., Lai, Y., Peng, M., Lopez-Lira, A., and Huang, J. (2023). Pixiu: A large language model, instruction data and evaluation benchmark for finance.

Xiong, F., Zhang, X., Feng, A., Sun, S., and You, C. (2025). QuantAgent: Price-Driven Multi-Agent LLMs for High-Frequency Trading. *arXiv preprint arXiv:2509.06784*.

Xu, W., Liang, Z., Mei, K., Gao, H., Tan, J., and Zhang, Y. (2025). A-Mem: Agentic memory for LLM agents. *arXiv preprint arXiv:2502.12110*.

Yang, H., Liu, X.-Y., and Wang, C. D. (2023). FinGPT: Open-source financial large language models.

Yoon, C., Lee, T., Hwang, H., Jeong, M., and Kang, J. (2024). CompAct: Compressing retrieved documents actively for question answering. In *EMNLP 2024*, pp. 21424–21439.

Yu, Y., Li, H., Chen, Z., Jiang, Y., Li, Y., Zhang, D., Liu, R., Suchow, J. W., and Khashanah, K. (2023). FinMem: A performance-enhanced LLM trading agent with layered memory and character design.

Zhang, B., Yang, H., and Liu, X.-Y. (2023a). Instruct-FinGPT: Financial sentiment analysis by instruction tuning of general-purpose large language models.

Zhang, H., Hua, F., Xu, C., Kong, H., Zuo, R., and Guo, J. (2024a). Unveiling the potential of sentiment: Can large language models predict Chinese stock price movements?

Zhang, W., Zhao, L., Xia, H., Sun, S., Sun, J., Qin, M., Li, X., Zhao, Y., Zhao, Y., Cai, X., Zheng, L., Wang, X., and An, B. (2024b). A multimodal foundation agent for financial trading: Tool-augmented, diversified, and generalist.

Zhang, Z., Bo, X., Ma, C., Li, R., Chen, X., Dai, Q., Zhu, J., Dong, Z., and Wen, J.-R. (2024c). A survey on the memory mechanism of large language model based agents. *arXiv preprint arXiv:2404.13501*.

Zheng, Y., Fu, D., Hu, X., Cai, X., Ye, L., Lu, P., and Liu, P. (2025). DeepResearcher: Scaling deep research via reinforcement learning in real-world environments. *arXiv preprint arXiv:2504.03160*.

Zhong, W., Guo, L., Gao, Q., Ye, H., and Wang, Y. (2024). MemoryBank: Enhancing large language models with long-term memory. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 38:19724–19731.

A Prompts

A.1 MACD Agent

To instantiate the **MACDAgent**, we construct a prompt that guides the agent to identify short-term momentum shifts using the **Moving Average Convergence Divergence (MACD)** indicator. Operating under high-frequency trading (HFT) constraints, the agent follows a two-step reasoning process: first generating the MACD chart via tool invocation, then analyzing key momentum patterns such as crossovers, histogram variation, and divergence before issuing a trading signal.

Prompt

You are a MACD (Moving Average Convergence Divergence) analysis assistant operating in a high-frequency trading context.
You must first call the tool `generate_macd_image` using the provided `kline_data`.
Once the MACD chart is generated, analyze the image for the following key signals:

1. **Signal Line Crossovers**: MACD line crossing above/below the signal line
2. **Zero Line Crossovers**: MACD line crossing above/below the zero line
3. **Histogram Analysis**: Changes in histogram direction and magnitude
4. **Divergences**: Price vs. MACD divergences (bullish/bearish)
5. **Momentum Strength**: Distance from zero line and slope steepness

Only after generating and analyzing the MACD image should you make a prediction about momentum direction and trading signals.
Do not make any predictions before generating and analyzing the image.

Prompt for MACDAgent in our multi-agent LLM framework. The agent receives OHLC data as input and first invokes `generate_macd_image` to obtain a visual depiction of market momentum. This prompt enforces a tool-first reasoning loop, ensuring predictions are grounded on the generated chart rather than raw prices.

Prompt

This is a —time_frame— MACD chart generated from recent OHLC market data. The chart includes:
MACD Line (blue) – difference between 12-period and 26-period EMAs;
Signal Line (red) – 9-period EMA of MACD line;
Histogram (bars) – difference between MACD and Signal lines;
Zero Line – reference for bullish/bearish momentum.
Analyze the chart with respect to the following aspects:

1. **Signal Line Crossovers:** Identify whether the MACD line crosses above (bullish) or below (bearish) the signal line and its relation to the zero line.
2. **Zero Line Crossovers:** Determine overall momentum and recent transitions across the zero baseline.
3. **Histogram Dynamics:** Evaluate bar expansion or contraction to gauge momentum strength.
4. **Divergence Patterns:** Compare price action with MACD behavior for bullish or bearish divergences.
5. **Momentum Strength:** Assess slope direction and distance from zero to estimate trend intensity.

Provide:

- **Primary Signal:** Buy, Sell, or Hold
- **Momentum Direction:** Bullish, Bearish, or Neutral
- **Confidence Level:** High, Medium, or Low
- **Supporting Evidence:** Observed MACD behavior
- **Risk Considerations:** Potential false or conflicting signals

Support your prediction with reasoning and observed evidence.

Graph-based prompt for MACDAgent in our multi-agent LLM framework. Given a generated MACD chart, the agent examines crossovers, histogram movements, and divergence to infer directional bias and confidence. This visual grounding enables consistent momentum interpretation across varying market conditions.

A.2 AVWAP Agent

To instantiate the **AVWAPAgent**, we construct a prompt that guides the agent to analyze price behavior relative to the **Anchored Volume Weighted Average Price (AVWAP)**. Operating under high-frequency trading (HFT) constraints, the agent combines tool-based chart generation with structured reasoning to detect support/resistance reactions, mean-reversion signals, and breakout dynamics around volume-weighted fair value.

Prompt

You are an Anchored VWAP (Volume Weighted Average Price) analysis assistant operating in a high-frequency trading context. You must first call the tool `generate_anchored_vwap_image` using the provided `kline_data`.
Once the Anchored VWAP chart is generated, analyze the image for the following key patterns:

1. **Support/Resistance Levels**: How price reacts when touching the AVWAP line.
2. **Mean Reversion Signals**: Extreme deviations from AVWAP suggesting pullbacks.
3. **Breakout Patterns**: Decisive moves above/below AVWAP with volume confirmation.
4. **AVWAP Slope Analysis**: Direction and steepness indicating market bias.
5. **Volume Confirmation**: Volume spikes during AVWAP interactions.
6. **Distance Analysis**: Magnitude of price deviation from fair value anchor.

Only after generating and analyzing the Anchored VWAP image should you make predictions about price direction and trading opportunities.
Do not make any predictions before generating and analyzing the image.

Prompt for AVWAPAgent in our multi-agent LLM framework. The agent first generates an Anchored VWAP chart from OHLC data, enforcing a visual-first reasoning flow. This ensures that the model grounds its interpretation on price–volume structure rather than raw numerical sequences.

Prompt

This is a —time_frame— Anchored VWAP chart generated from recent OHLC market data. The chart shows the volume-weighted average price anchored to a significant market event or time point.
Analyze the chart with respect to the following aspects:

1. **Support/Resistance Analysis:** Determine whether AVWAP acts as support (price above) or resistance (price below) and assess reactions upon contact.
2. **Mean Reversion Opportunities:** Evaluate deviations from AVWAP and identify potential pullbacks toward fair value.
3. **Breakout/Breakdown Signals:** Detect decisive moves above or below AVWAP, especially with volume confirmation.
4. **AVWAP Slope and Bias:** Interpret the slope's direction (upward = bullish, downward = bearish) and steepness to infer bias strength.
5. **Fair Value and Volume Context:** Assess whether price trades at a premium or discount to AVWAP and how volume supports current price action.

Provide:

- **Primary Trade Signal:** Buy, Sell, or Hold with rationale
- **AVWAP Bias:** Bullish, Bearish, or Neutral
- **Fair Value Assessment:** Premium, Discount, or Fair
- **Entry Strategy:** Suggested positioning based on current setup
- **Risk Management:** Stop-loss and profit target zones

- **Volume Confirmation:** How volume aligns with price behavior

Support your prediction with reasoning and observed signals.

Graph-based prompt for AVWAPAgent in our multi-agent LLM framework. The agent interprets AVWAP as a dynamic reference of institutional fair value, analyzing slope, distance, and volume interplay to infer bias and entry timing. This visual grounding allows robust short-term reasoning around market equilibrium and mean-reversion behavior.

A.3 RSI Agent

To instantiate the **RSIAgent**, we construct a prompt that enables the agent to evaluate momentum strength and reversal potential using the **Relative Strength Index (RSI)** indicator. Operating under high-frequency trading (HFT) constraints, the agent first generates an RSI chart via tool invocation, then interprets overbought/oversold conditions, divergences, and swing structures to infer directional bias.

Prompt

You are an RSI (Relative Strength Index) analysis assistant operating in a high-frequency trading context.
You must first call the tool `generate_rsi_image` using the provided `kline_data`.
Once the RSI chart is generated, analyze the image for the following key signals:

1. **Overbought/Oversold Levels**: RSI crossing above 70 or below 30.
2. **Centerline Crossovers**: RSI crossing above/below 50.
3. **Failure Swings**: RSI breaking prior swing points after overbought/oversold moves.
4. **Divergences**: Price–RSI mismatches (bullish/bearish).
5. **Hidden Divergences**: Trend-continuation signals during pullbacks.
6. **RSI Trendlines**: RSI breakout patterns anticipating price breakouts.
7. **Swing Rejections**: RSI rejection patterns near extreme zones.

Only after generating and analyzing the RSI image should you make a prediction about momentum direction and trading signals.
Do not make any predictions before generating and analyzing the image.

Prompt for RSIAgent in our multi-agent LLM framework. The agent first produces an RSI chart from OHLC input, enforcing a tool-first reasoning step. This ensures that its predictions are visually grounded in oscillator behavior rather than inferred directly from raw data.

Prompt

This is a —time_frame— RSI chart generated from recent OHLC market data. The plot shows:
RSI Line (0–100 scale), **Overbought Level** (70), **Oversold Level** (30), and **Center Line** (50).
Analyze the chart with respect to the following aspects:

1. **Overbought/Oversold Conditions:** Identify sustained or recent entries into extreme zones and crossovers signaling reversals.
2. **Centerline Crossovers:** Determine whether RSI is above or below 50 and assess momentum bias.
3. **Failure Swings:** Detect RSI reversals breaking prior highs/lows after overbought or oversold readings.
4. **Divergences:** Compare price and RSI trajectories for bullish or bearish divergences, including hidden ones.
5. **Trendlines and Patterns:** Check for RSI trendline or pattern breakouts that may precede price shifts.
6. **Swing Rejections:** Evaluate bullish or bearish rejection sequences within overbought/oversold regions.

Provide:

- **Primary Signal:** Buy, Sell, or Hold

- **Momentum Direction:** Bullish, Bearish, or Neutral
- **Confidence Level:** High, Medium, or Low
- **Supporting Evidence:** Observed RSI behavior
- **Risk Considerations:** Potential false or conflicting signals

Support your prediction with reasoning and observed patterns.

Graph-based prompt for RSIAgent in our multi-agent LLM framework. Given the generated RSI chart, the agent assesses momentum strength, reversals, and divergences to infer directional bias. This visually grounded reasoning improves short-term momentum interpretation and reversal detection consistency across different market conditions.

A.4 Heiken Ashi Agent

To instantiate the **HeikenAshiAgent**, we construct a prompt that guides the agent to identify directional trends and momentum transitions using **Heiken Ashi (HA)** candlestick representations. Operating under high-frequency trading (HFT) conditions, the agent first generates a Heiken Ashi chart and then interprets trend strength, reversals, and momentum consistency to support short-term directional reasoning.

Prompt

You are a Heiken Ashi candlestick analysis assistant operating in a high-frequency trading context.
You must first call the tool `generate_heiken_ashi_image` using the provided `kline_data`.
Once the Heiken Ashi chart is generated, analyze the image for the following key patterns:

1. **Trend Direction**: Consecutive green (bullish) or red (bearish) candles.
2. **Trend Strength**: Candle body size, wick formation, and consistency.
3. **Reversal Signals**: Color changes, doji formations, or wicks indicating indecision.
4. **Momentum Analysis**: Body-size acceleration or deceleration.
5. **Consolidation Patterns**: Mixed colors and small bodies showing indecision.
6. **Entry/Exit Timing**: Ideal trade points derived from HA formations.

Only after generating and analyzing the Heiken Ashi image should you make predictions about trend direction and trading opportunities.
Do not make any predictions before generating and analyzing the image.

Prompt for HeikenAshiAgent in our multi-agent LLM framework. The agent receives OHLC input and invokes `generate_heiken_ashi_image` to visualize smoothed candlestick dynamics. This tool-first approach ensures the reasoning process is grounded in visually observable trend patterns rather than raw market data.

Prompt

This is a —time_frame— Heiken Ashi chart generated from recent OHLC data. The chart smooths price action to reduce noise and highlight trend direction.
Green candles represent bullish momentum, **red candles** represent bearish momentum, and wicks indicate indecision or volatility.
Analyze the chart according to the following aspects:

1. **Trend Direction & Consistency:** Determine whether consecutive candles form a clear bullish or bearish sequence and detect color transitions signaling possible reversals.
2. **Trend Strength:** Assess candle body size, wick presence, and overall uniformity to gauge momentum.
3. **Reversal Indications:** Identify doji-like or spinning-top formations, color shifts after long trends, or weakening candle bodies.
4. **Momentum Dynamics:** Observe

whether candle bodies expand (acceleration) or contract (deceleration).

5. **Consolidation vs Trending:** Detect mixed-color candles and narrow ranges indicating indecision.

6. **Entry/Exit Signals:** Evaluate pullback entries or exhaustion exits based on Heiken Ashi structure.

Provide:

- **Trend Status:** Strong Bullish, Weak Bullish, Strong Bearish, Weak Bearish, or Sideways
- **Momentum Phase:** Accelerating, Steady, Decelerating, or Reversing
- **Signal Quality:** High, Medium, or Low
- **Trading Action:** Buy, Sell, Hold, or Wait
- **Risk Assessment:** Probability of false or conflicting signals

Support your prediction with reasoning and evidence observed from the chart.

Graph-based prompt for HeikenAshiAgent in our multi-agent LLM framework. Given the generated Heiken Ashi chart, the agent identifies dominant color sequences, momentum acceleration, and potential reversal setups. This visual grounding enhances consistency in trend identification and trade-timing inference across high-frequency market conditions.

A.5 SMA Agent

To instantiate the **SMAAgent**, we construct a prompt that enables the agent to identify trend direction and momentum shifts using the **Simple Moving Average (SMA)** indicator. Operating under high-frequency trading (HFT) constraints, the agent first generates an SMA chart via tool invocation, then interprets crossovers, slope dynamics, and price alignment relative to multiple SMAs.

Prompt

You are a Simple Moving Average (SMA) analysis assistant operating in a high-frequency trading context.
You must first call the tool `generate_sma_image` using the provided `kline_data`.
Once the SMA chart is generated, analyze the image for the following key signals:

1. **Trend Identification**: Price position relative to key SMAs (above/below).
2. **Moving Average Crossovers**: Faster SMA crossing above/below slower SMA.
3. **Dynamic Support/Resistance**: Price interacting with SMA levels.
4. **Multiple MA Ribbon**: Alignment and spacing between multiple SMAs.
5. **Mean Reversion**: Distance of price from SMA suggesting potential reversal.
6. **SMA Slope**: Direction and steepness indicating trend strength.

Only after generating and analyzing the SMA image should you make predictions about trend direction and trading signals.
Do not make any predictions before generating and analyzing the image.

Prompt for SMAAgent in our multi-agent LLM framework. The agent receives OHLC data and generates an SMA chart as a visual reference for trend detection. This approach ensures that reasoning about market direction is grounded in moving-average structure rather than isolated price movements.

Prompt

This is a —time_frame— SMA chart generated from recent OHLC market data. The chart includes multiple Simple Moving Averages (e.g., 20-, 50-, and 200-period), each capturing a different time horizon of trend behavior.
Analyze the chart with respect to the following aspects:

1. **Trend Identification:** Assess whether price is above or below key SMAs and infer the overall directional bias.
2. **Crossovers:** Identify golden or death crosses between short and long SMAs as trend reversal signals.
3. **Support/Resistance:** Evaluate SMA lev-

els where price consistently bounces or reverses.

4. **MA Ribbon Alignment:** Determine if SMAs are neatly ordered (strong trend) or entangled (ranging phase).
5. **Mean Reversion:** Identify overextended price movements far from SMAs suggesting pullbacks.
6. **Slope and Momentum:** Examine the slope of each SMA to assess momentum direction and strength.

Provide:

- **Primary Trend Direction:** Bullish, Bearish, or Range-bound
- **Trend Strength:** Strong, Moderate, or Weak
- **Trading Bias:** Buy on Dips, Sell on Rallies, or Stay Flat
- **Key SMA Levels:** Notable support/resistance zones
- **Confirmation Needed:** Other indicators that would validate the trend

Support your prediction with structured reasoning and observed SMA configurations.

Graph-based prompt for SMAAgent in our multi-agent LLM framework. Given the generated SMA chart, the agent interprets crossover events, slope direction, and ribbon spacing to estimate trend intensity and bias. This visual reasoning approach improves short-term adaptability while maintaining consistency in trend assessment.

A.6 Stochastic Agent

To instantiate the **StochasticAgent**, we construct a prompt that enables the agent to analyze momentum reversals and exhaustion levels using the **Stochastic Oscillator**. Operating under high-frequency trading (HFT) conditions, the agent first generates a Stochastic chart, then interprets crossovers, divergences, and overbought/oversold patterns to infer short-term momentum bias.

Prompt

You are a Stochastic Oscillator analysis assistant operating in a high-frequency trading context.
You must first call the tool `generate_stochastic_image` using the provided `kline_data`.
Once the Stochastic chart is generated, analyze the image for the following key signals:

1. **Overbought/Oversold Conditions**: Levels above 80 (overbought) or below 20 (oversold).
2. **Line Crossovers**:
3. **Divergences**: Price–Stochastic mismatches (bullish/bearish).
4. **Failure Swings**: Internal divergence patterns signaling reversals.
5. **Centerline Crossovers**: Crossings above/below the 50 level.
6. **Slope Analysis**: Direction and steepness of

Only after generating and analyzing the Stochastic image should you make predictions about momentum direction and trading signals. Do not make any predictions before generating and analyzing the image.

Prompt for StochasticAgent in our multi-agent LLM framework. The agent begins by invoking a visualization tool to produce the Stochastic chart, ensuring that subsequent reasoning about reversals and momentum bias is grounded in oscillator behavior rather than abstract data interpretation.

Prompt

This is a —time_frame— Stochastic Oscillator chart generated from recent OHLC data. It plots the **%K line** (fast) and the **%D line** (slow) along with overbought (80), oversold (20), and center (50) reference levels.
Analyze the chart according to the following aspects:

1. **Overbought/Oversold Conditions:** Identify whether the oscillator is above 80 or below 20 and if it is exiting these

zones.

2. **Line Crossovers:** Determine where the %K crosses the %D, especially within extreme regions.

3. **Divergences:** Compare price and Stochastic behavior to detect bullish or bearish divergences.

4. **Failure Swings:** Identify internal Stochastic divergences signaling momentum exhaustion.

5. **Centerline Position:** Assess whether the oscillator is above or below 50 to determine directional bias.

6. **Slope and Momentum:** Evaluate the steepness and direction of %K and %D for acceleration or deceleration.

Provide:

- **Primary Signal:** Buy, Sell, or Hold
- **Momentum Direction:** Bullish, Bearish, or Neutral
- **Confidence Level:** High, Medium, or Low
- **Supporting Evidence:** Key Stochastic patterns observed
- **Risk Considerations:** Potential false or conflicting signals

Support your prediction with structured reasoning and explicit pattern-based justification.

Graph-based prompt for StochasticAgent in our multi-agent LLM framework. Once the Stochastic chart is generated, the agent evaluates oscillator dynamics to detect reversals, divergences, and short-term exhaustion points. This process allows visually anchored reasoning for momentum-based entry and exit signals in fast-moving markets.

A.7 Trend Agent

To instantiate the **TrendAgent**, we construct a prompt that directs the agent to identify short-term directional bias based on **support and resistance trendlines** derived from recent K-line data. Operating under high-frequency trading (HFT) conditions, the agent first generates a trendline chart and then interprets slope, compression, and breakout behavior to infer near-term market direction.

Prompt

You are a K-line trend pattern recognition assistant operating in a high-frequency trading context.
You must first call the tool `generate_trend_image` using the provided `kline_data`.
Once the chart is generated, analyze the image for support/resistance trendlines and notable candlestick formations.
Only after generating and analyzing the image should you make a prediction about the short-term trend (upward, downward, or sideways).
Do not make any predictions before generating and analyzing the image.

Prompt for TrendAgent in our multi-agent LLM framework. The agent receives recent OHLC data and first visualizes it with automated trendlines. This ensures that all reasoning about directional bias and breakout probability is grounded in geometric structure and price interaction rather than raw numerical features.

Prompt

This is a —time_frame— K-line chart with automatically generated trendlines: the **blue line** represents support, and the **red line** represents resistance, both derived from recent closing prices.
Analyze the chart based on the following aspects:

1. **Trendline Interaction:** Observe whether price is bouncing off, breaking through, or consolidating between trendlines.

2. **Slope Direction:** Identify whether the dominant slope is upward (bullish), downward (bearish), or flat (neutral).

3. **Compression/Expansion:** Detect narrowing or widening distances between trendlines as signs of volatility contraction or breakout setup.

4. **Candlestick Behavior:** Assess wick direction, candle clustering, and breakout attempts along support/resistance zones.

5. **Breakout Probability:** Estimate the likelihood of a trendline breach based on slope angle and recent price clustering.

Provide:

- **Predicted Trend:** Upward, Downward, or Sideways
- **Trend Strength:** Strong, Moderate, or Weak
- **Key Evidence:** Trendline behavior and candlestick interactions
- **Risk Considerations:** False breakout potential or low-volatility traps

Support your prediction with structured reasoning and clear trendline-based evidence.

Graph-based prompt for TrendAgent in our multi-agent LLM framework. Once the chart is generated, the agent analyzes trendline slope, spacing, and price compression to infer directional momentum. This visual grounding enables consistent short-term trend interpretation across volatile high-frequency market states.

A.8 Decision Agent

To instantiate the **DecisionAgent**, we construct two prompts around the final action step: a pre-decision synthesis that fuses all agent reports with retrieved memories, and a post-decision reflection that audits the outcome for learning and calibration. This preserves the tool-first, evidence-grounded style used across META while adapting it to a decision-and-reflect workflow.

Prompt

You are an elite HFT quantitative analyst making real-time trade decisions on **{stock name}** using the **{time frame}** timeframe.
MISSION: Issue an immediate execution order **(LONG or SHORT only; HOLD is prohibited)** based on comprehensive analysis.
FORECAST HORIZON: Predict market direction for the next 1–2 candlesticks ({time frame} period).
HISTORICAL CONTEXT — Similar Market Conditions: If similar memories are provided, summarize key lessons and outcomes; otherwise state that no highly relevant memories were found.
CURRENT MARKET ANALYSIS (provided in context): Indicator, Pattern, Trend, MACD, Anchored VWAP, Heiken Ashi, SMA, Stochastic, and RSI reports.
DECISION FRAMEWORK

- **Primary Signals (High Weight):** Momentum confirmations (e.g., MACD cross, RSI extremes), completed breakouts with volume, decisive trendline breaks.
- **Secondary Signals (Medium Weight):** Multi-indicator alignment, S/R interactions, near-complete patterns.
- **Risk Filters:** Downweight conflicting signals unless one side dominates; reflect historical outcomes; set **risk–reward** in **[1.2, 1.8]** based on volatility and conviction.
- **Memory Integration:** Reference comparable past scenarios, outcomes, and confidence adjustments.

REQUIRED OUTPUT FORMAT (JSON): Return only valid JSON

```
{
 "forecast_horizon":"Next {time_frame}
  candlestick prediction",
  "decision":"LONG or SHORT",
 "confidence_level":"HIGH/MEDIUM/LOW",
  "primary_drivers":["2-3 strongest
  signals"],
  "memory_insights":"Key lessons from
  similar scenarios",
  "justification":"Concise reasoning
  combining current analysis with
  history",
  "risk_reward_ratio":"<float between
  1.2 and 1.8>",
  "stop_loss_rationale":"Brief risk
  management rationale"
}
```

Decision Hierarchy

1. **Strong Alignment:** All reports + positive memory outcomes → High confidence trade.
2. **Moderate Alignment:** 2/3 reports align + supportive memory → Medium confidence.

3. **Weak Signals:** Mixed reports → Choose direction with strongest recent momentum and best precedent.

Execute with precision. The market waits for no one.

Prompt for DecisionAgent in our multi-agent LLM framework. The agent fuses multi-indicator evidence with memory context and outputs a **single** LONG/SHORT decision in strict JSON, including confidence and risk–reward calibration.

Prompt

You are conducting a post-trade analysis as an elite HFT quantitative analyst.
TRADE SUMMARY (given): Stock, Timeframe, Decision, Confidence, Expected Risk–Reward.
ORIGINAL ANALYSIS SUMMARY (given): Concise record of drivers, memory insights, justification, and risk management.
ACTUAL TRADE OUTCOME (given): P&L, final status (profitable/loss/breakeven), and realized market behavior.
REFLECTION FRAMEWORK

1. **Decision Quality:** Was the analysis sound? Did primary drivers anticipate the move? Did memory insights help?
2. **Execution:** Interpretation fidelity of indicators/patterns; appropriateness of confidence.
3. **Learning:** Under/overestimated conditions; most/least reliable signals; how to adapt next time.
4. **Memory Integration:** How does this outcome update memory; any weight shifts for indicators.

REQUIRED OUTPUT: Produce a concise 3–5 sentence reflection covering: key factors behind the outcome; what worked/failed; concrete lessons; any adjustments to approach or confidence calibration. Focus on actionable improvements for similar future setups.

Post-decision reflection prompt for DecisionAgent in our multi-agent LLM framework. The reflection converts outcome feedback into memory-aligned guidance, tightening future signal weighting and confidence calibration.

B Computation and Mathematical Formulation of Indicators

This appendix provides the formal definitions and computational procedures for the six technical indicators integrated within the META (Memory-Enhanced Trading Agent) architecture. Each indicator serves as a distinct analytical perspective on market dynamics—capturing momentum, mean reversion, trend persistence, or value anchoring. All indicators are computed from raw OHLCV (Open–High–Low–Close–Volume) data and act as the quantitative foundation for the multi-agent reasoning framework described in Appendix A.

B.1 Moving Average Convergence Divergence (MACD)

Definition. The Moving Average Convergence Divergence (MACD) is a momentum indicator derived from the difference between two exponential moving averages (EMAs) of price, designed to quantify acceleration and deceleration of market trends. Formally, for price series P_t:

$$\text{MACD}_t = \text{EMA}_{\text{fast},t} - \text{EMA}_{\text{slow},t},$$

where

$$\text{EMA}_{k,t} = \alpha_k P_t + (1-\alpha_k)\text{EMA}_{k,t-1}, \quad \alpha_k = \frac{2}{k+1}.$$

Common parameters are $(\text{fast}, \text{slow}) = (12, 26)$.

The signal line is defined as:

$$\text{Signal}_t = \text{EMA}_9(\text{MACD}_t),$$

and the histogram, which reflects short-term divergence between trend speeds, is:

$$\text{Hist}_t = \text{MACD}_t - \text{Signal}_t.$$

Interpretation. The MACD measures momentum by comparing short- and long-term EMAs. A positive value implies upward acceleration; negative values imply downward momentum. Crossover events between MACD and the signal line indicate potential reversals, while histogram expansion suggests strengthening momentum. Because MACD captures the rate of change in trend direction, it is frequently used in HFT environments for detecting early inflection points.

B.2 Relative Strength Index (RSI)

Definition. The Relative Strength Index (RSI) quantifies the internal strength of recent price movements by comparing average gains to average losses within a defined lookback period n (typically $n = 14$). For a price difference $\Delta P_t = P_t - P_{t-1}$:

$$\text{AvgGain}_t = \text{EMA}_n(\max(\Delta P_t, 0))$$

$$\text{AvgLoss}_t = \text{EMA}_n(\max(-\Delta P_t, 0)).$$

The relative strength (RS) and RSI are then computed as:

$$RS_t = \frac{\text{AvgGain}_t}{\text{AvgLoss}_t}, \qquad \text{RSI}_t = 100 - \frac{100}{1 + RS_t}.$$

Interpretation. RSI values oscillate between 0 and 100, serving as a bounded momentum oscillator. Readings above 70 denote overbought conditions (potential reversal risk), while readings below 30 suggest oversold conditions (potential recovery). Crossings of the 50 midpoint line are interpreted as transitions in trend dominance. In high-frequency contexts, RSI slope and failure-swing formations provide early warnings of momentum exhaustion before price reversal becomes visible.

B.3 Anchored Volume Weighted Average Price (AVWAP)

Definition. The Anchored Volume Weighted Average Price (AVWAP) extends the conventional VWAP by fixing the accumulation anchor to a specific event time τ (e.g., major news or breakout). It represents the cumulative volume-weighted mean of traded prices since the anchor point:

$$\text{AVWAP}_t = \frac{\sum_{i=\tau}^{t} P_i V_i}{\sum_{i=\tau}^{t} V_i},$$

where P_i denotes the typical or closing price and V_i the corresponding trading volume.

Interpretation. AVWAP serves as a dynamic "fair value" benchmark representing the mean price paid per traded volume since a specific anchor. When price trades above AVWAP, the market is said to be in a bullish premium phase; when below, in a bearish discount phase. Because institutional participants often reference anchored VWAP levels for liquidity decisions, these lines act as adaptive support or resistance zones in HFT systems.

B.4 Simple Moving Average (SMA)

Definition. The Simple Moving Average (SMA) is a non-weighted mean of the most recent n closing prices:

$$\text{SMA}_t = \frac{1}{n} \sum_{i=0}^{n-1} P_{t-i}.$$

For two SMAs with different windows $n_1 < n_2$, the crossover between SMA_{n_1} and SMA_{n_2} provides directional signals.

Interpretation. SMA acts as a linear filter that removes short-term noise to reveal the underlying trend. Shorter-period SMAs respond quickly to price fluctuations, while longer-period SMAs define macro-trend bias. The alignment of multiple SMAs into a monotonic order (short > medium > long) indicates trend strength, whereas frequent crossovers reflect low directional conviction or range-bound conditions.

B.5 Stochastic Oscillator

Definition. The Stochastic Oscillator measures the relative position of the current close within a recent price range of length n:

$$\%K_t = 100 \times \frac{P_t - L_n}{H_n - L_n}, \quad \%D_t = \text{SMA}_3(\%K_t),$$

where $H_n = \max(P_{t-n+1}, \ldots, P_t)$ and $L_n = \min(P_{t-n+1}, \ldots, P_t)$ denote the highest and lowest prices in the past n periods (commonly $n = 14$).

Interpretation. The Stochastic Oscillator quantifies momentum relative to the recent range rather than absolute changes. Readings above 80 signify potential overbought conditions, while values below 20 indicate oversold conditions. Crossovers between %K and %D provide early entry and exit signals. In HFT environments, the slope and curvature of %K serve as short-term acceleration indicators, and divergence between Stochastic and price action often precedes directional reversal.

B.6 Heiken Ashi (HA)

Definition. Heiken Ashi candlesticks are a transformation of traditional OHLC data that average and smooth price information to reduce noise. The recursive formulation is:

$$\text{HA_Close}_t = \tfrac{1}{4}(O_t + H_t + L_t + C_t),$$

$$\text{HA_Open}_t = \tfrac{1}{2}(\text{HA_Open}_{t-1} + \text{HA_Close}_{t-1}),$$

$$\text{HA_High}_t = \max(H_t, \text{HA_Open}_t, \text{HA_Close}_t),$$

$$\text{HA_Low}_t = \min(L_t, \text{HA_Open}_t, \text{HA_Close}_t).$$

Interpretation. Heiken Ashi averages smooth intra-period volatility to better capture directional persistence. Extended sequences of uniform-color candles represent strong trends, while small-bodied candles or alternating colors indicate weakening momentum or consolidation. This transformation is particularly useful in high-frequency applications, where standard candlestick noise can obscure subtle but statistically persistent directional biases.

C Case Study

This case study illustrates a full trading cycle within META, from signal perception to reflective memory update. Each trade begins with input *K-line* data analyzed in parallel by specialized **Indicator Agents** (e.g., MACD, RSI, AVWAP, SMA, Stochastic, Heiken Ashi), each generating a structured signal report. These outputs form an **indicator vector** that encodes the current market state and is used to retrieve historically similar episodes from the memory base through cosine similarity search. The retrieved memories, together with real-time analyses, are then passed to the **Decision Agent** for synthesis into a trading action (LONG, SHORT, or HOLD) with associated confidence. Following trade execution, the **Reflection Module** evaluates the outcome and updates the episodic memory—adding, refining, or pruning trade instances to ensure continual learning and adaptive decision-making.

Trend Agent Output Response

1. Trendline Interaction:

- The price is interacting with both the support (blue) and resistance (red) lines.
- Recent candles have bounced off the support line, indicating it is holding for now.
- The resistance line has been tested multiple times, with the price failing to break through.

2. Trendline Slope:

- Both the support and resistance lines are sloping downward, indicating a bearish trend.

3. Recent K-line Behavior:

- The price is compressing between the support and resistance lines, showing consolidation.
- There is a slight downward bias as the price is closer to the support line.

Prediction: Downward Trend
Reasoning:

- The downward slope of both trendlines suggests a bearish trend.
- The price is closer to the support line, and repeated tests of this line may lead to a breakdown.
- The inability to break above the resistance line indicates strong selling pressure.

Signals:

- Watch for a break below the support line for confirmation of a downward move.
- If the price breaks above the resistance line, it could indicate a reversal, but current signals favor a downward trend.

MACD Agent Output Response

MACD Analysis:
1. Signal Line Crossovers:

- The MACD line (blue) is currently below the signal line (red), indicating a bearish crossover.
- This crossover occurs below the zero line, reinforcing bearish momentum.

2. Zero Line Crossovers:

- The MACD line is below zero, indicating negative momentum.
- No recent crossovers of the zero line, suggesting sustained bearish momentum.

3. Histogram Analysis:

- The histogram bars are contracting, indicating weakening bearish momentum.
- No recent histogram zero-line crossovers, but the contraction suggests a potential momentum shift.

4. Divergence Patterns:

- No clear bullish or bearish divergences observed between price action and MACD behavior.
- Price lows align with MACD lows, confirming current trend.

5. Momentum Strength:

- The MACD line is significantly below zero, indicating strong bearish momentum.
- Both MACD and signal lines are sloping downward, but the slope is decreasing, suggesting deceleration.

Conclusion:

- **Primary Signal:** Hold
- **Momentum Direction:** Bearish
- **Confidence Level:** Medium
- **Key Supporting Evidence:** Bearish crossover below zero line, contracting histogram.
- **Risk Considerations:** Potential for false signals due to histogram contraction; watch for any bullish crossover or divergence for a trend reversal.

AVWAP Agent Output Response

1. Support/Resistance Analysis:

- **Price Position:** The price is currently below the AVWAP line.
- **Price Reaction:** The price has been rejected at the AVWAP line multiple times, indicating it is acting as resistance.
- **Dynamic Role:** AVWAP is acting as dynamic resistance in a downtrend.

2. Mean Reversion Opportunities:

- **Deviation:** The price has deviated significantly below the AVWAP, suggesting a potential pullback to fair value.
- **Extreme Stretches:** The current stretch below AVWAP suggests a possible mean reversion opportunity.
- **Typical Range:** The price has shown a tendency to revert to the AVWAP after deviations.

3. Breakout/Breakdown Signals:

- **Decisive Breaks:** There have been decisive breaks below the AVWAP with strong volume.
- **Reclaiming/Losing AVWAP:** The price has lost the AVWAP after being above it.
- **Volume Behavior:** Volume spikes coincide with AVWAP tests and breakdowns.

4. AVWAP Slope and Bias:

- **Slope:** The AVWAP line is sloping downward, indicating a bearish bias.
- **Steepness:** The slope is moderately steep, suggesting a strong directional bias.
- **Slope Dynamics:** The slope is consistent, maintaining a bearish trend.

5. Fair Value Assessment:

- **Institutional Fair Value:** The AVWAP level represents institutional fair value.
- **Price Position:** Current prices are trading at a discount to volume-weighted fair value.
- **Evolution:** The fair value relationship has shifted from premium to discount as the session progressed.

6. Volume-Price Relationship:

- **Volume Increase:** Volume increases when the price approaches the AVWAP.
- **Volume Spikes:** Notable volume spikes occur during AVWAP tests and breakdowns.
- **Volume Support:** Volume supports the current bearish price action relative to AVWAP.

Comprehensive Assessment:

- **Primary Trade Signal: Sell**
  - **Rationale:** The price is below a downward-sloping AVWAP, acting as resistance with strong volume confirmation.
- **AVWAP Bias: Bearish**
  - **Based on:** Downward slope and price position below AVWAP.
- **Fair Value Assessment: Discount**
  - **Relative to AVWAP:** Current prices are below the volume-weighted fair value.
- **Entry Strategy:**
  - **Approach:** Enter short positions on rallies towards the AVWAP with confirmation of rejection.
- **Risk Management:**
  - **Stop Loss:** Above recent highs near the AVWAP.
  - **Profit Target:** At recent lows or further below, considering trend continuation.
- **Volume Confirmation:**
  - **Support:** Volume supports the bearish signal, with spikes during AVWAP interactions confirming resistance.

This analysis suggests a continuation of the downtrend, with AVWAP serving as a key resistance level.

Heiken Ashi Agent Output Response

Heiken Ashi Analysis
1. Trend Direction & Consistency:

- **Consecutive Candles:** There is a sequence of green candles followed by a longer sequence of red candles.
- **Dominant Color:** Red candles dominate, indicating a bearish trend.
- **Recent Color Changes:** A few green candles appear at the end, suggesting a potential shift.
- **Trend Cleanliness:** The trend is relatively clean with minimal alternation during the bearish phase.

2. Trend Strength Indicators:

- **Candle Bodies:** Initially large green bodies, followed by large red bodies, indicating strong momentum in both directions.
- **Wicks:** Red candles have minimal upper wicks, showing strong bearish momentum.
- **Body Size:** Decreasing body size towards the end of the red sequence suggests deceleration.

3. Reversal Pattern Recognition:

- **Doji-like Formations:** Some small-bodied candles with wicks appear, indicating indecision.
- **Color Change:** Green candles after a strong red sequence suggest a potential reversal.
- **Spinning Tops/Hammer-like Patterns:** Present towards the end of the red sequence.

4. Momentum Acceleration/Deceleration:

- **Candle Size:** Recent candles are smaller, indicating fading momentum.
- **Gaps:** No significant gaps observed.
- **Wick Reduction:** Progressive wick reduction in the bearish direction.

5. Consolidation vs Trending Phases:

- **Color Alternation:** Recent alternation suggests a potential consolidation phase.
- **Small Bodies with Long Wicks:** Indicate range-bound action.

6. Entry/Exit Signal Quality:

- **Signal Quality:** Medium confidence due to recent indecision and potential reversal signals.
- **Optimal Entry Points:** Consider buying if a strong green candle confirms a reversal.
- **Trend Exhaustion:** Watch for further green candles to confirm trend change.

7. Multi-Candle Pattern Analysis:

- **Three Consecutive Same-Color Candles:** Present in both green and red sequences, confirming trends.
- **Two-Candle Reversal Patterns:** Potential reversal pattern forming with recent green candles.
- **Extended Trend Sequences:** Red sequence nearing exhaustion.

Comprehensive Assessment:

- **Trend Status:** Weak Bearish transitioning to Sideways
- **Momentum Phase:** Decelerating
- **Signal Quality:** Medium
- **Trading Action:** Wait for better setup
- **Entry Strategy:** Look for a strong green candle to confirm reversal before entering a buy position.
- **Exit Criteria:** Exit if red candles reappear with strong bodies and minimal wicks.
- **Risk Assessment:** Moderate probability of false signals due to recent indecision and small-bodied candles.

SMA Agent Output Response

SMA Analysis:
1. Trend Identification:

- **Price/SMA Relationship:** The price is currently below both the 20-period (blue) and 40-period (red) SMAs, indicating a bearish trend.
- **SMA Alignment:** The 20-period SMA is below the 40-period SMA, reinforcing the downtrend.

2. Moving Average Crossovers:

- **Death Cross:** The 20-period SMA has crossed below the 40-period SMA, confirming a bearish signal.
- **Recent Crossovers:** No recent golden cross observed; the last significant crossover was bearish.

3. Dynamic Support/Resistance:

- **Support:** Price is not finding support at the SMAs; it is below both.
- **Resistance:** The 20-period SMA is acting as resistance during rallies.
- **Respect of Levels:** Price has consistently stayed below the SMAs, respecting them as resistance.

4. Multiple MA Ribbon Analysis:

- **SMA Stacking:** The SMAs are neatly stacked with the shorter below the longer, indicating a strong bearish trend.
- **Ribbon Behavior:** The ribbon is fanning out, suggesting a strengthening downtrend.

5. Mean Reversion Signals:

- **Price Extension:** Price is significantly extended below the SMAs, indicating potential for a mean reversion or pullback.
- **Deviation:** The price is stretched, suggesting a possible short-term reversal.

6. Slope and Alignment:

- **SMA Slopes:** Both SMAs are sloping downwards, indicating a strong bearish trend.
- **Alignment:** SMAs are aligned in the same downward direction, confirming trend strength.
- **Slope Steepness:** The slopes are moderately steep, indicating a strong trend.

Conclusion:

- **Primary Trend Direction:** Bearish
- **Trend Strength:** Strong
- **Trading Bias:** Sell on Rallies
- **Key SMA Levels:** Watch the 20-period SMA ($\sim$86000) as resistance.
- **Confirmation Needed:** Look for confirmation from other indicators like RSI for oversold conditions or MACD for momentum confirmation.

This analysis suggests a strong bearish trend with potential short-term pullbacks. Traders should consider selling on rallies, especially if price approaches the 20-period SMA resistance.

Stochastic Oscillator Agent Output Response

Stochastic Oscillator Analysis:
1. Overbought/Oversold Conditions:

- The Stochastic is currently below the 20 level, indicating oversold conditions.
- It has been in the oversold zone for a short period.
- There is no clear exit from the oversold zone yet, suggesting a potential continuation of the current trend.

2. Line Crossovers:

- The %K line is crossing above the %D line, which is a bullish signal.
- This crossover is occurring in the oversold zone, which can indicate a potential reversal.

3. Divergence Patterns:

- There is no clear bullish divergence as the price and Stochastic both show lower lows.
- No bearish divergence is present.

4. Failure Swings:

- No clear bullish failure swing is observed as the Stochastic has not made a higher low in the oversold territory.
- No bearish failure swing is present.

5. Centerline Analysis:

- The Stochastic is below the 50 centerline, indicating bearish momentum.
- A crossover above 50 would suggest a shift to bullish momentum, but this has not occurred yet.

6. Slope and Momentum:

- The %K line is starting to slope upwards, indicating a potential increase in bullish momentum.
- The %D line is also beginning to turn upwards, supporting the potential for a bullish reversal.
- Both lines are making lower highs and lows, consistent with the current downtrend.

Conclusion:

- **Primary Signal:** Hold
- **Momentum Direction:** Neutral to Bullish (potential reversal)
- **Confidence Level:** Medium
- **Key Supporting Evidence:** Bullish crossover in oversold zone, upward slope of %K and %D lines.
- **Risk Considerations:** The Stochastic is still in the oversold zone, and the overall trend is bearish. There is a risk of false signals if the price does not confirm the reversal.

This analysis suggests caution, as the bullish signals are emerging but not yet confirmed by price action.

RSI Agent Output Response

RSI Analysis:
1. Overbought/Oversold Conditions:

- The RSI is not currently in overbought (>70) or oversold (<30) territory. It is around the 40 level.
- There is no recent crossover from these extreme zones.

2. Centerline Crossovers:

- The RSI is below the centerline (50), indicating bearish momentum.
- There have been no recent crossovers of the 50 level.

3. Failure Swings:

- No clear failure swings are observed as RSI hasn't reached extreme levels recently.

4. Divergence Patterns:

- **Bullish Divergence:** The price made a lower low, while RSI made a higher low, suggesting potential bullish reversal.

5. RSI Trendlines & Patterns:

- No visible trendlines or chart patterns like triangles are apparent on the RSI.

6. Swing Rejection Patterns:

- No clear swing rejection patterns are observed.

7. RSI Range Analysis:

- RSI is ranging between 30 and 50, showing bearish momentum but not extreme exhaustion.
- Recent support around 30 and resistance near 50.

Conclusion:

- **Primary Signal:** Hold
- **Momentum Direction:** Bearish
- **Confidence Level:** Medium
- **Key Supporting Evidence:** Bullish divergence observed, but overall momentum remains bearish.
- **Risk Considerations:** Potential for false signals due to lack of strong reversal patterns and RSI not reaching extreme levels.

The bullish divergence suggests a possible reversal, but the overall bearish momentum and lack of strong reversal signals warrant caution.

After analyzing the current market state, META encodes the aggregated indicator outputs into an *indicator vector* and retrieves the most similar past trades from memory. These analogous cases serve

as experiential references, anchoring the decision process in prior market behaviors.

Retrieved Memory 1 (Similarity: 0.89)
Date: 2025-09-30T17:47:07.801365
Forecast Horizon: Next 4-hour candlestick prediction
Confidence Level: Medium
Primary Drivers:

- Descending Triangle pattern indicating potential bearish breakout
- Price below VWAP and SMAs, confirming bearish trend

Memory Insights: No highly relevant past trade memories found for current market conditions.
Justification: The analysis suggested a bearish sentiment with indicators pointing toward oversold conditions. The descending triangle and price positioning below key moving averages supported a continuation of the downward trend.
Risk–Reward Ratio: 1.5
Stop-Loss Rationale: Stop loss placed above recent resistance to mitigate reversal risk in alignment with the bearish setup.
Decision: SHORT
Outcome: 387.31
Reflection: The trade's success stemmed from correctly identifying the descending triangle and confirming bearish momentum via VWAP and SMA alignment. Effective risk management further supported profitability. However, the medium confidence rating highlights an opportunity to enhance robustness through expanded historical referencing and refined indicator weighting in future decisions.

Retrieved Memory 2 (Similarity: 0.83)
Date: 2025-09-17T15:16:46.572820
Forecast Horizon: Next 4-hour candlestick prediction
Confidence Level: Medium
Primary Drivers:

- Price below both 20 and 40-period SMAs indicating a bearish trend
- MACD line significantly below signal line with negative histogram indicating strong bearish momentum
- Price below downward-sloping AVWAP acting as resistance

Memory Insights: Previous scenarios with similar bearish setups have shown mixed outcomes, highlighting the need for cautious confidence and dynamic risk management.
Justification: Technical indicators including the MACD, SMAs, and AVWAP suggested a continuation of the bearish trend. Historical analysis supported this decision, though with mixed results.
Risk–Reward Ratio: 1.5
Stop-Loss Rationale: Stop loss placed above recent highs near the AVWAP to manage reversal risk.
Decision: SHORT
Outcome: –76.69
Reflection: The loss arose from overreliance on technical bearish signals that failed to materialize. Future improvements include better weighting of historical insight, adaptive stop-loss adjustment, and accounting for broader market factors.

Retrieved Memory 3 (Similarity: 0.83)
Date: 2025-09-18T16:40:44.859425
Forecast Horizon: Next 4-hour candlestick prediction
Confidence Level: Medium
Primary Drivers:

- Price below both 20 and 40-period SMAs indicating a bearish trend
- MACD line below signal line with negative histogram indicating bearish momentum
- Descending Triangle pattern suggesting bearish breakout

Memory Insights: Similar setups with descending triangles and bearish indicators have produced mixed outcomes, calling for cautious confidence.
Justification: The bearish alignment of MACD, SMAs, and chart pattern supported a short position. Historical results validated this setup.
Risk–Reward Ratio: 1.5
Stop-Loss Rationale: Stop loss placed above recent highs near the AVWAP.
Decision: SHORT
Outcome: 571.08
Reflection: Profit resulted from precise identification of the descending triangle and indicator alignment. The trade confirms the importance of weighted indicator calibration and dynamic confidence scaling.

Retrieved Memory 4 (Similarity: 0.82)
Date: 2025-09-16T17:21:41.742237
Forecast Horizon: Next 4-hour candlestick prediction
Confidence Level: Medium
Primary Drivers:

- MACD line below signal line with negative histogram indicating bearish momentum
- Price below both 20 and 40-period SMAs confirming downtrend
- Descending Triangle pattern suggesting potential breakdown

Memory Insights: Prior successful shorts in similar descending triangle setups with bearish indicators.
Justification: Strong bearish signals supported a short trade. Historical parallels suggested profit potential.
Risk–Reward Ratio: 1.5
Stop-Loss Rationale: Stop loss placed above recent highs near the AVWAP.
Decision: SHORT
Outcome: –150.00
Reflection: The loss resulted from overestimating bearish momentum. Future improvements include multi-signal confirmation, contextual confidence scaling, and indicator reweighting.

Retrieved Memory 5 (Similarity: 0.81)
Date: 2025-09-20T00:33:33.390160
Forecast Horizon: Next 4-hour candlestick prediction
Confidence Level: Medium
Primary Drivers:

- RSI and Stochastic Oscillator in oversold territory indicating potential reversal
- Bullish crossover in Stochastic Oscillator
- Price trading below VWAP suggesting potential undervaluation

Memory Insights: Oversold conditions and bullish crossovers often lead to short-term reversals in past cases.
Justification: Despite a bearish backdrop, RSI, Stochastic, and VWAP indicated undervaluation warranting a cautious long.
Risk–Reward Ratio: 1.5
Stop-Loss Rationale: Stop loss placed below recent lows to mitigate continuation risk.
Decision: LONG
Outcome: 76.69
Reflection: Success stemmed from recognizing oversold conditions and aligning multiple reversal signals. Future enhancements include dynamic indicator weighting and adaptive confidence tuning based on confluence strength.

The retrieved memories provide valuable contextual grounding for the current trade. By examining both profitable and unprofitable precedents, META's synthesis stage integrates lessons from historical performance into the present decision-making process. Profitable cases such as Memory 3 highlight the effectiveness of aligning descending triangle patterns with bearish indicator convergence, reinforcing confidence in similar market setups. Conversely, losses observed in Memory 2 and Memory 4 underscore the risks of overreliance on technical signals without sufficient contextual validation or adaptive stop-loss control. Through reflective comparison, META dynamically adjusts the weighting of key indicators—such as MACD, SMA alignment, and AVWAP slope—while moderating confidence when past patterns exhibit mixed outcomes. This learning mechanism enables the agent to refine its strategy beyond static rule-following, improving robustness by encoding both the strengths and weaknesses of prior trades into future reasoning.

Decision Agent Output Response
Forecast Horizon: Next 4h candlestick prediction
Decision: SHORT
Confidence Level: Medium
Primary Drivers:

- Descending Triangle pattern indicating potential bearish breakout
- Price below both 20 and 40-period SMAs confirming downtrend
- MACD line below signal line with negative histogram indicating bearish momentum

Memory Insights: Previous similar setups with descending triangle patterns and price below key moving averages have resulted in mixed outcomes, highlighting the need for cautious confidence and dynamic risk management.
Justification: The current market analysis shows a strong bearish trend with the price below key moving averages and a descending triangle pattern suggesting a potential breakdown. Historical memory supports the decision to short in similar conditions, where the alignment of technical indicators led to both successful and unsuccessful outcomes. The overall bearish sentiment is reinforced by the price's position below key moving averages and the AVWAP.
Risk–Reward Ratio: 1.5
Stop-Loss Rationale: Place stop loss above recent highs near the AVWAP to manage risk in case of a reversal.

Simulated Trade Outcome
Outcome: WIN
PnL: +615.30
Entry Price: 82,548.81
Exit Price: 81,933.51
Performance Summary: The trade achieved a profitable outcome, with price movement aligning closely to the anticipated bearish continuation. Execution adhered to the risk–reward target of 1.5, validating both directional accuracy and the stop-loss placement strategy.

Trade Reflection
The key factors that led to the profitable outcome were the accurate identification of the descending triangle pattern and the confirmation of bearish momentum through the MACD and moving averages. The original analysis was sound, as the primary drivers correctly predicted market movement, and historical memory insights provided a balanced perspective on risk management. However, the medium confidence level suggests there was room for greater conviction given the alignment of technical indicators. For future scenarios, it would be beneficial to emphasize pattern strength and multi-indicator confluence, adjusting confidence upward in similar conditions. Additionally, refining stop-loss placement to account for volatility could further improve risk management and trade robustness.

Following the trade's completion and reflection phase, META transitions into the final stage of its cognitive workflow—the **Memory Consolidation Phase**. In this stage, the *Memory Operator* (Ψ) evaluates the outcome of the trade alongside retrieved precedents to determine how the memory base should evolve. Based on the reflection, past memories may be updated, retained, or removed, while the newly completed trade is encoded as a

new episodic memory. This process ensures that the agent's experiential knowledge remains both adaptive and performance-oriented, continually refining its future decision quality.

Memory Agent Output Response
Post-Trade Memory Operations:

- **Memory 1:** NOOP — retained without modification.
- **Memory 2:** NOOP — maintained as historical reference.
- **Memory 3:** UPDATE — performance data appended with revised confidence weighting.
- **Memory 4:** DELETE — removed due to repeated underperformance in similar contexts.
- **Memory 5:** UPDATE — outcome consistency improved; metadata refined for future retrieval weighting.
- **New Memory:** ADD — stored as a new episodic record representing the current trade outcome and reflection.

Summary: The Memory Operator integrates experiential feedback by updating profitable analogs, pruning uninformative cases, and adding the newly completed trade as a fresh memory entry. This targeted consolidation ensures the memory base remains relevant, adaptive, and performance-aligned across future retrieval cycles.

This case study illustrates META's complete cognitive workflow, demonstrating how the system transforms perception-driven analysis into experience-informed reasoning. By integrating insights from multiple indicator agents, META constructs a structured representation of the market state that captures both short-term and long-term dynamics. The retrieval of analogous trade memories provides contextual grounding, enabling the decision agent to synthesize new judgments with awareness of past outcomes. Finally, the reflection and memory consolidation stages close the cognitive loop, allowing the agent to refine its internal knowledge base through targeted updates and selective forgetting.

Through this process, META exhibits several key advantages: (1) it enhances predictive stability by grounding real-time reasoning in historical evidence; (2) it maintains interpretability through explicit agent-level outputs; and (3) it achieves continual learning without retraining, instead relying on reflection-driven memory management. Together, these mechanisms position META as a step toward genuinely adaptive, self-improving trading intelligence capable of evolving alongside dynamic market conditions.

D Metrics Design

To ensure comparability with prior work, we adopt the evaluation metrics introduced in QUANTAGENT (**?**), which quantify both predictive accuracy and profitability under risk-constrained execution. All metrics are computed within a fixed three-candlestick horizon, corresponding to the short-term forecasting window defined by the trading agent.

Directional Accuracy (α). Directional accuracy measures the model's consistency in predicting the correct price movement over the next three candlesticks. For a LONG decision, each candle that closes above the current close counts as a correct prediction; for a SHORT decision, each candle that closes below the current close counts. The accuracy is given by:

$$\alpha = \frac{C}{T},$$

where C denotes the number of correctly predicted candles and $T = 3$ the evaluation horizon. This formulation aligns with the Mean Directional Accuracy metric used in financial forecasting (**?**).

Rate-of-Return (RoR) Metrics. To evaluate trade profitability and risk exposure, we measure multiple forms of rate of return between the entry and exit prices, incorporating realistic stop-loss and take-profit execution. The base stop-loss threshold ρ is fixed at 0.0005 (0.05%), while the reward threshold is determined by the model-generated risk–reward ratio r, yielding a maximum allowed gain $R = r \cdot \rho$ and maximum loss $-r \cdot \rho$:

$$R = r \cdot \rho, \quad \rho = 0.0005.$$

Each trade is exited at the first intra-candle price that hits either threshold.

- **Cumulative Return (R_{cc}):** The realized return over the forecast horizon, reflecting the net gain or loss achieved at the earliest hit of either the take-profit or stop-loss limit:

$$R_{cc} = \begin{cases} \dfrac{P_{\text{exit}} - P_{\text{entry}}}{P_{\text{entry}}}, & \text{if LONG,} \\ \dfrac{P_{\text{entry}} - P_{\text{exit}}}{P_{\text{entry}}}, & \text{if SHORT.} \end{cases}$$

  This simulates a risk-managed intra-candle execution with bounded return outcomes.

- **Maximum Return ($R_{\max}$):** The best-case rate of return achievable over the same window, assuming an optimal intra-candle exit:

$$R_{\max} = \begin{cases} \dfrac{P_{\text{high}} - P_{\text{entry}}}{P_{\text{entry}}}, & \text{if LONG,} \\ \dfrac{P_{\text{entry}} - P_{\text{low}}}{P_{\text{entry}}}, & \text{if SHORT.} \end{cases}$$

  $R_{\max}$ represents the upper bound of potential profit.

- **Minimum Return ($R_{\min}$):** The most adverse price movement within the same horizon:

$$R_{\min} = \begin{cases} \dfrac{P_{\text{low}} - P_{\text{entry}}}{P_{\text{entry}}}, & \text{if LONG,} \\ \dfrac{P_{\text{entry}} - P_{\text{high}}}{P_{\text{entry}}}, & \text{if SHORT.} \end{cases}$$

  $R_{\min}$ captures the worst-case drawdown under identical trade direction.

- **Similarity-weighted Return Alignment (R_{sim}):** This metric assesses the consistency between similarity scores of retrieved memories and realized profitability. For each trade i, with similarity weights w_{ij} over k retrieved episodes and their respective realized returns $R_{cc}^{(j)}$, we define:

$$R_{\text{sim}} = \frac{1}{N} \sum_{i=1}^{N} \frac{\sum_{j=1}^{k} w_{ij} \cdot R_{cc}^{(j)}}{\sum_{j=1}^{k} w_{ij}},$$

  where N is the number of evaluated trades. Higher R_{sim} values indicate that the retrieval mechanism successfully identifies memories whose return patterns align with the current market context.

Together, these metrics provide a comprehensive view of agent performance: α measures directional accuracy; R_{cc} reflects realized profitability; $R_{\max}$ and $R_{\min}$ establish upper and lower risk bounds; and R_{sim} quantifies the semantic consistency between retrieved experiences and market outcomes. This unified metric suite captures predictive precision, profitability, and memory relevance under realistic HFT execution conditions, fully aligned with the QuantAgent evaluation protocol.

E Use of Large Language Models (LLMs)

Large Language Models (LLMs) were employed solely as supportive instruments to improve the clarity, coherence, and grammatical precision of the manuscript. Specifically, GPT-5 was utilized to refine portions of the text—primarily within the Introduction and Methodology sections—to ensure readability and stylistic consistency. The authors retain full responsibility for all scientific and technical contributions, including the formulation of research questions, experimental design, implementation of the META framework, and validation of empirical findings. The use of LLMs was limited strictly to editorial assistance and did not influence the conceptual, analytical, or experimental aspects of this work.