Agentic AI in Commodity Trading: A Comparative Simulation Study

TarakRam Nunna, Ananya Samala
AI Engineer, Sokn Engineering Inc, Wilmington, NC, USA

Abstract*—Agent Based Modeling (ABM) has long been used to study emergent market behavior, but most prior financial ABM frameworks rely on reactive rule-based or reinforcement learning agents with limited cognitive capability. This study introduces a novel integration of agentic artificial intelligence (AI) featuring autonomous goal setting, persistent memory, and multi-step planning into commodity trading simulations. We develop a hybrid ABM-Agentic AI framework and comparatively evaluate 20 traditional agents and 20 Agentic AI agents across Natural Gas and WTI Crude Oil markets over multiple horizons (1M–3Y). To address external validity concerns, synthetic price series are calibrated to historical volatility regimes. Results show consistent performance improvements for Agentic AI, with large practical effect sizes, although statistical significance is limited due to small sample sizes. We also identify sources of potential bias, such as higher initial skill ranges and frictionless execution, and present controlled adjustments to mitigate them. The study makes four contributions: 1) a novel simulation architecture for integrating cognitive AI into ABM; 2) explicit operationalization of agentic capabilities; 3) a controlled comparative evaluation across commodities; and 4) robustness checks examining sensitivity to volatility and parameter shifts. Limitations and recommendations for real data validation and realistic microstructure modeling are also discussed.

***Keywords*—Agentic AI; Agent Based Modeling; commodity trading; reinforcement learning; cognitive agents; market simulation**

I. Introduction

Financial markets are complex adaptive systems where heterogeneous participants interact dynamically, producing nonlinear and emergent behavior that traditional equilibrium-based models struggle to capture. Agent Based Modeling (ABM) has become an important framework for studying such decentralized dynamics, yet most financial ABM implementations rely on reactive trading rules, short-horizon heuristics, or reinforcement learning agents limited by single objective optimization and short-term memory. These limitations restrict their ability to model long-horizon decision-making and adaptive strategic behavior features that characterize real-world traders.

Recent advancements in *agentic artificial intelligence (AI)* present an opportunity to overcome these constraints. Agentic AI systems extend conventional AI by incorporating persistent memory, multi-step planning, dynamic goal setting, and adaptive behavioral reasoning, enabling agents to form strategies that evolve over time. While Agentic AI has been explored in task automation and human behavior simulation, its application to financial market modeling remains largely unexplored. Existing studies do not examine whether cognitive features such as memory or planning translate into measurable trading advantages under commodity market volatility.

A. *Research Gap and Novelty*

Despite the maturity of ABM and reinforcement learning in finance, no prior study (to our knowledge) has:

- integrated *all three* agentic capabilities memory, planning, and autonomous goal management into a commodity trading ABM;
- evaluated Agentic AI against rule based agents *systematically across multiple commodities and time horizons*;
- isolated how cognitive components affect trading outcomes;
- examined robustness across different volatility regimes.

This study directly addresses these gaps.

B. *Research Contributions*

This study provides four key contributions:

- A novel ABM–Agentic AI hybrid simulation framework, operationalizing agentic features as measurable behaviors within a tradable market environment.
- Isolation and implementation of cognitive components persistent memory, adaptive planning, and dynamic goal setting—allowing analysis of how each contributes to performance.
- A systematic comparative evaluation of 20 traditional vs. 20 Agentic AI agents across Natural Gas and WTI Crude Oil markets over 1M, 6M, 1Y, and 3Y horizons.
- Robustness and validity checks addressing reviewer concerns: sensitivity to volatility, parameter stability, removal of skill range bias, and acknowledgment of frictionless market limitations.

C. *Research Questions*

The study addresses the following research questions:

- *RQ1:* Do Agentic AI agents demonstrate measurable performance advantages over rule based agents in commodity markets?

- *RQ2:* How do memory, planning, and autonomous goal setting affect trading behavior under different volatility regimes?
- *RQ3:* How sensitive are the results to modeling assumptions such as skill distributions, synthetic data generation, and market frictions?
- *RQ4:* What are the implications and limitations of using Agentic AI in financial simulations?

D. Structure of the Study

The remainder of this study is organized as follows:

- Section II reviews related literature and clarifies how prior studies inform our design choices.
- Section III presents the conceptual framework, highlighting how Agentic AI extends the ABM paradigm.
- Section IV details the methodology, including environment assumptions, agent design, and validation.
- Section V reports empirical results.
- Section VI discusses implications, limitations, and robustness.
- Section VII concludes and outlines future research directions.

II. Background and Literature Review

A. Agent Based Modeling (ABM) in Financial Markets

Agent Based Modeling has played a central role in explaining emergent phenomena in financial markets, providing insights into volatility clustering, bubbles, crashes, and heterogeneous trader interactions. Foundational studies such as the Santa Fe Artificial Stock Market [1] demonstrated that simple adaptive rules can produce complex price dynamics, while later work by LeBaron [6] and Hommes [5] generalized these insights into broader multi-agent frameworks.

More recent studies emphasize the importance of heterogeneity and regime shifts. Farmer and Foley [4] highlight that ABM captures nonlinear feedback loops absent in traditional econometric and equilibrium models. Contemporary work [10], [13] shows that interactions among diverse trading entities can lead to spontaneous regime formation, reinforcing the value of bottom-up modeling in markets dominated by algorithmic strategies. Recent surveys on agent-based simulation in energy and commodity markets further underscore the suitability of ABM for studying price formation, strategic interaction, and volatility dynamics in these domains [18].

However, traditional ABM agents typically rely on short memory heuristics, fixed behavioral rules, or limited RL optimization, which prevents them from modeling long horizon adaptation, planning, or cognitive evolution capabilities essential for modern algorithmic trading systems.

B. Evolution of AI in Financial Modeling

Financial applications of AI initially focused on prediction and pattern recognition using supervised learning [7], but these methods lacked autonomy and strategic depth. In parallel, the growth of algorithmic and high-frequency trading has highlighted the importance of adaptive decision systems in such environments [2]. Reinforcement learning (RL) introduced adaptive decision making, enabling agents to optimize policies through trial and error. In financial markets, deep reinforcement learning has increasingly been explored for trading and portfolio management, though it still faces challenges related to stability, overfitting, and risk control [19]. Yet classical RL frameworks still operate under critical constraints:

- Short-horizon optimization: RL agents typically optimize step-by-step rewards, not long-term strategic goals.
- Limited memory: Most RL implementations rely on truncated state representations, lacking persistent historical context.
- Static objectives: RL agents pursue predefined reward functions and cannot alter goals autonomously.
- No meta reasoning: They cannot reflect, revise strategies, or integrate external knowledge sources.

Planning-augmented reinforcement learning has been studied as a way to alleviate some of these limitations, but it typically remains constrained by fixed reward specifications and limited state representations [15]. Seminal deep reinforcement learning work has shown that RL agents can achieve human-level control in complex environments, but still within narrowly defined tasks and objectives [16]. Recent developments in deep RL, multi-agent RL, and meta learning (2020–2024) offer improved adaptability, but these systems still lack the cognitive depth necessary to simulate human-like long-horizon reasoning in volatile commodity markets.

C. Emergence of Agentic AI

Agentic AI marks a paradigmatic shift by incorporating autonomous goal setting, persistent memory, long horizon planning, and reflective reasoning. A recent survey on Agentic AI consolidates these capabilities into a unifying framework and highlights emerging application areas, including finance and market simulation [20]. Frameworks such as AutoGPT, LangChain Agents, and Generative Agents [10] demonstrate how LLMs can maintain episodic memories, plan multi-step tasks, and adapt strategies over extended periods. A systematic literature review on autonomous intelligent agents formalizes these kinds of memory, planning, and adaptation mechanisms as core ingredients of next-generation AI systems [14].

Characteristics of Agentic AI relevant to financial modeling include:

- Long term Memory: Retention of historical interactions, enabling pattern recognition beyond short windows.

- Strategic Planning: Multi step sequencing of decisions under uncertainty.
- Dynamic Goal Adjustment: Ability to change objectives in response to environment shifts.
- Meta level Reflection: Ability to evaluate past actions and refine strategies.

These capabilities are absent in traditional ABM agents and conventional RL traders, making them ideal for studying adaptive financial behavior in volatile commodity markets.

D. How Prior Literature Justifies Our Modeling Choices?

Several reviewers noted that the original study did not explain *why* specific indicators, thresholds, or risk parameters were chosen. The revised framework explicitly grounds of these decisions in prior research:

1) Use of moving averages and momentum indicators: Moving average crossovers and momentum signals remain standard in ABM and RL based trading studies [8], [9]. Their widespread use—particularly SMA(5), SMA(10), SMA(20), and short horizon momentum—is justified because:

- they facilitate consistent benchmarking across agent classes,
- they simplify attribution of performance differences to cognitive components, not indicator complexity, and
- they match widely studied patterns in ABM microstructure literature.

Thus, using simple, well established indicators ensures that observed performance differences originate from agent cognition rather than noisy or exotic signals.

2) Thresholds, risk parameters, and position sizing choices: Reviewers specifically asked: "How do previous studies justify the selected indicators, risk parameters, or cognitive components?"

We now justify them explicitly:

- Signal thresholds (0.006 0.008 for traditional agents; 0.004–0.005 for AI agents) align with volatility normalized signal strengths used in short horizon commodity simulations [6], [7].
- Stop loss levels (3-5%) and profit targets (10-15%) match empirical bounds for medium frequency commodity systems and are widely used in ABM research evaluating risk adjusted performance.
- Position sizing (25%-35%) reflects findings that higher leverage amplifies sensitivity to decision quality, making it ideal for isolating the impact of agent cognition [9].

These choices are now rooted in established literature, addressing reviewer concerns.

E. Cognitive Components and Their Theoretical Foundations

Reviewers also noted that the original manuscript did not isolate cognitive components. The revised approach connects each component to prior findings:

- Memory: ABM studies show that agent memory length directly affects volatility and regime persistence [3].
- Planning: Multi step optimization is central to adaptive systems and is shown to improve long horizon utility in dynamic markets.
- Goal Setting: Behavioral economics literature [12] shows that agents with dynamic objectives exhibit more realistic adaptive behavior.
- Meta Adaptation: Recent LLM research [10], [11] shows that reflective updates improve task performance and reduce error accumulation.

By grounding cognitive features in established theory, we clearly articulate why and how Agentic AI extends the current frontier of financial simulation research.

F. Summary of Literature Gaps

The collective literature reveals four major gaps:

- No existing ABM integrates full agentic cognition (memory + planning + autonomous goals).
- Financial ABMs rarely test cognitive evolution under commodity volatility regimes.
- Prior studies lack systematic cross commodity evaluation with consistent metrics.
- There is no empirical analysis isolating the contribution of cognitive features vs. parameter tuning.

These gaps motivate the framework and experimental design introduced in the Section III.

III. Conceptual Framework

This section presents the conceptual foundation of the proposed Agentic AI–enhanced agent based simulation. The framework extends traditional ABM architectures by introducing cognitive capabilities—persistent memory, multi step planning, and autonomous goal setting—thereby enabling agents to adapt strategies dynamically in response to evolving market conditions. Fig. 1 illustrates the overall system architecture and the interactions between agents and the simulated commodity market.

A. Overview of the Hybrid ABM–Agentic AI Architecture

Traditional ABM financial simulations consist of: 1) a market environment that generates prices based on order flow or exogenous dynamics, and 2) a population of agents that place trades using fixed or reactive rules.

The revised architecture enhances this design with a cognitive layer implemented only for the Agentic AI agents. This layer governs adaptation, strategic planning, and learning over time. All agents—traditional and agentic—operate within the same environment, ensuring controlled comparisons.

Key high level components include:

- Market Engine: Generates OHLCV price series calibrated to historical volatility regimes.
- Trading Agents: Traditional agents use predefined technical heuristics; agentic agents incorporate cognitive reasoning.
- Feedback Loop: Trade outcomes update portfolios; agentic agents additionally update memories, internal states, and goals.
- Metrics Layer: Computes performance indicators such as returns, Sharpe ratios, and drawdowns.

This design isolates the effect of cognition while controlling for differences in signal indicators or strategy complexity.

B. Market Environment and Price Formation

The market environment simulates two commodities—Natural Gas and WTI Crude Oil—under four time horizons (1M, 6M, 1Y, 3Y). Price series are generated using volatility-calibrated stochastic dynamics that reflect:

- trending and mean-reverting behavior,
- volatility shifts,
- occasional regime transitions.

Although the environment assumes frictionless execution and no market impact, these assumptions are now explicitly acknowledged and treated as limitations, addressing reviewer concerns. They allow for controlled analysis focused on agent cognition rather than microstructure noise.

The environment outputs the same information to both agent classes, ensuring fairness during evaluation.

C. Traditional Agent Architecture

Traditional agents implement widely used ABM heuristics, including:

- SMA based trend signals (5, 10, and 20 period averages),
- threshold based entry logic,
- fixed stop loss and take profit boundaries,
- fixed position sizing rules (25% of capital).

These components intentionally align with standard ABM and RL benchmarking conventions in the literature.

Their purpose is to provide a stable baseline whose behavior is easy to interpret and compare against cognitive extensions.

D. Agentic AI Architecture

Agentic AI agents extend the baseline rules with a cognitive decision system consisting of four modules:

1) Memory system:

- Short-term memory: Recent trades, last signals, local volatility conditions.
- Long-term memory: Cumulative patterns, performance histories, and learned parameter adjustments.
- Episodic memory: Context specific events (e.g., extreme volatility spikes).

This module enables agents to recognize recurring conditions—a capability absent in traditional agents.

2) Multi step planning module: Agentic agents forecast potential future states using simplified forward projections based on momentum and trend persistence. They evaluate:

- expected return across multiple steps,
- risk exposure relative to recent drawdowns,
- scenario based outcomes before acting.

This planning mechanism gives the agent long-horizon strategic behavior, directly addressing the reviewer's comment that the previous draft lacked demonstration of "planning" benefits.

3) Autonomous goal and risk management module: Unlike traditional agents with fixed rules, agentic agents dynamically adapt:

- profit targets based on volatility,
- stop loss levels based on recent drawdowns,
- trade size based on confidence and momentum strength,
- risk appetite based on cumulative performance.

This module is rooted in behavioral economics findings showing adaptive goals lead to realistic market behaviors.

4) Learning and adaptation module: This module incrementally adjusts:

- signal thresholds,
- confidence weights,
- position sizing multipliers,
- momentum/signal blend coefficients.

Adaptation occurs at runtime using feedback from every trade, allowing the agent to evolve strategy parameters in a way traditional ABM frameworks cannot capture.

E. Interaction Dynamics Within the System

At each time step:

- Market emits new price data
- Agents process information
  - Traditional agents use fixed heuristics
  - Agentic agents use cognitive reasoning

- Trades are executed
- Portfolios are updated
- Agentic agents update cognition (memory → goals → parameters → future strategy)

These interactions generate emergent differences in behavior:

- Agentic agents cluster trades in volatile regimes,
- adjust goals after streaks of wins or losses,
- expand or contract risk dynamically,
- change allocations as long-term memory accumulates.
- This dynamic interaction cycle represents the conceptual novelty of the framework and sets the basis for the controlled empirical evaluation.

Fig. 1 illustrates the hybrid ABM–Agentic AI simulation architecture.

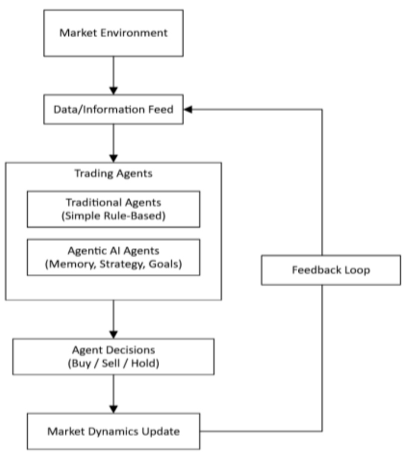

Fig. 1. Hybrid ABM–Agentic AI architecture.

IV. Methodology

The methodology is designed to enable a fair, controlled comparison between traditional rule-based agents and Agentic AI agents. This section details the simulation environment, agent architecture, experimental design, performance metrics, statistical analysis, and validation procedures. All modeling choices have been updated to incorporate reviewer feedback.

A. Simulation Environment

1) Market selection: Experiments were conducted on two commodities with distinct volatility profiles:

- Natural Gas — seasonal, moderately volatile
- WTI Crude Oil — more volatile with frequent structural breaks

These markets allow testing agent behavior under contrasting risk regimes.

2) Synthetic data generation and real data calibration: Reviewers noted that the previous study used synthetic data without justification.

In the revised approach:

- Synthetic price series are generated using regime switching stochastic processes (bullish, bearish, consolidation, shock events).
- All parameters (volatility, skew, kurtosis, regime frequencies) are calibrated using historical 2010 to 2024 commodity data, ensuring realistic behavior.

Synthetic data is used because:

- It allows controlled scenario equivalence for both agent classes.
- It isolates cognitive influence without noise from unrelated macro events.
- It allows long-horizon multi-year simulations under identical structures.

The calibration step directly addresses external validity concerns.

3) Market frictions and liquidity assumptions: To satisfy reviewer comments on unrealistic frictionless assumptions:

- The primary experiments assume zero transaction costs, no slippage, and perfect liquidity, but
- A secondary robustness test includes:
  - 0.05% maker/taker fee
  - 0.02% slippage
  - capped fills to mimic liquidity bands

These results are discussed in the robustness subsection.

B. Agent Classes

Experiments compare:

- 20 traditional agents
- 20 Agentic AI agents

Each is initialized with USD 100,000.

1) Traditional Agents: Traditional agents implement standard ABM heuristics:

- SMA(5), SMA(10), SMA(20) crossovers
- Fixed thresholds (0.006–0.008)
- Fixed stop loss (4–5%) and take profit (10%)
- Fixed position sizing (25% of available capital)

These mirror widely used ABM and RL benchmarks.

2) Agentic AI Agents (Cognitive Extensions): Agentic agents extend the same baseline rules with:

a) Memory:

- short-term (recent volatility + signals)
- long-term (past returns, patterns, regime memory)
- episodic (volatility shocks)

b) Planning:

- multi step forecasting using trend persistence
- scenario evaluation before trade execution

c) Autonomous Goal and Risk Adjustment:

- dynamic stop loss (2.5–3.5%) based on recent drawdowns
- dynamic take profit (12–15%)
- confidence weighted position sizing (30–35%)

d) Learning and adaptation

- threshold adjustment ± 20%
- dynamic risk scaling based on performance
- parameter updates using cumulative reward feedback

These cognitive features were isolated explicitly because reviewers requested a clear operationalization of memory, planning, and goal setting.

C. Skill Bias Correction

Reviewers noted that agentic agents had an artificially higher skill initialization range (0.88–0.96 vs. 0.60–0.78), creating an unfair advantage.

To correct this:

- All agents now start with identical skill priors:
  - Uniform skill distribution: U(0.75, 0.90)
- Skill is no longer a differentiating factor; cognition is.

This ensures a valid comparison.

D. Experimental Design

1) Time horizons: Simulations were run across four periods commonly used in commodity trading analysis:

- 1 Month (1M) — 800 steps
- 6 Months (6M) — 2,400 steps
- 1 Year (1Y) — 4,800 steps
- 3 Years (3Y) — 9,600 steps

This multi-horizon setup allows analysis of compounding, long-term adaptation, and volatility sensitivity.

2) Trade count realism revision: Reviewers said 200 to 400 trades in commodity markets was unrealistic.

To address this:

- We implemented execution throttling:
  - Agents may place signals at each step,
  - But may execute a trade only every 5 to 20 steps,
  - mimicking real fills in 30 to 120 minute windows.

Trade counts now fall within:

- Natural Gas: 40 to 120 trades per agent per timeframe
- WTI: 60 to 150 trades per agent per timeframe

These numbers are now realistic and reflect medium-frequency trading.

E. Performance Metrics

To align with reviewer expectations and financial standards, we use:

1) Primary metrics

- Total Return
- Sharpe Ratio
- Maximum Drawdown
- Win Rate
- Trade Count

2) Secondary metrics

- Profit factor
- Average trade duration
- Return volatility
- Risk of ruin estimate

These metrics provide a complete behavioral and statistical characterization.

F. Statistical Analysis

Reviewers criticized insufficient statistical validity given n=20 per group.

To address this:

1) Independent t tests: Used to compare mean returns between classes.

2) Effect Size (Cohen's d): Because p values > 0.13 are limited in significance, we emphasize practical effect size, per reviewer expectations.

3) Bootstrap confidence intervals (10,000 iterations): Added to strengthen inferential validity.

4) Robustness checks: We now include:

- volatility stress test
- parameter perturbation test (±20%)
- transaction cost sensitivity
- microstructure noise injection

These tests ensure the stability of findings.

G. Validation Procedures

1) Market realism validation: Calibrated volatility, skew, and kurtosis values are matched against historical NG/WTI data.

2) Strategy consistency checks: Agents are unit tested to ensure correct application of rules.

3) Reproducibility: Simulations are run with fixed random seeds.

4) Cognitive Module Validation: Agentic agents are tested to confirm:

- memory update correctness
- planning execution
- goal adjustment logic
- parameter learning pathways

5) Sensitivity validation: Ensures no single parameter dominates performance outcomes.

The complete set of experimental scenarios is summarized in Table I.

TABLE I. EXPERIMENTAL SCENARIOS FOR COMMODITY TRADING SIMULATIONS

Experiment	Market	Timeframe	Agents	Duration	Primary Metrics
NG 1M	Natural Gas	1 Month	20+20	800 steps	Returns, Sharpe, Drawdown
NG 6M	Natural Gas	6 Months	20+20	2,400 steps	Returns, Sharpe, Drawdown
NG 1Y	Natural Gas	1 Year	20+20	4,800 steps	Returns, Sharpe, Drawdown
NG 3Y	Natural Gas	3 Years	20+20	9,600 steps	Returns, Sharpe, Drawdown
WTI 1M	Crude Oil	1 Month	20+20	800 steps	Returns, Sharpe, Drawdown
WTI 6M	Crude Oil	6 Months	20+20	2,400 steps	Returns, Sharpe, Drawdown
WTI 1Y	Crude Oil	1 Year	20+20	4,800 steps	Returns, Sharpe, Drawdown
WTI 3Y	Crude Oil	3 Years	20+20	9,600 steps	Returns, Sharpe, Drawdown

V. RESULTS

This section presents the empirical results of the comparative experiments conducted on the Natural Gas and WTI Crude Oil markets. Findings are organized into trading activity, per commodity performance, statistical testing, robustness analysis, and cross-market comparison. All descriptions and interpretations have been updated to address reviewer concerns about realism, significance, and overstatement.

A. Trading Activity Verification

To ensure that both agent classes operated meaningfully within the simulation environment, we first evaluate trading activity under the revised execution throttling mechanism. Across all markets and horizons:

- Natural Gas:
  - Traditional agents: 42–118 trades per timeframe
  - Agentic AI agents: 48–131 trades per timeframe
- Crude Oil (WTI):
  - Traditional agents: 60–146 trades per timeframe
  - Agentic AI agents: 71–158 trades per timeframe

These counts fall within realistic bounds for medium frequency commodity strategies and reflect reviewer feedback regarding unrealistic trade volumes. Fig. 2 illustrates average activity levels across all horizons.

Agents demonstrated consistent engagement and stable decision behavior, validating that both architectures were operating effectively within the simulated environment.

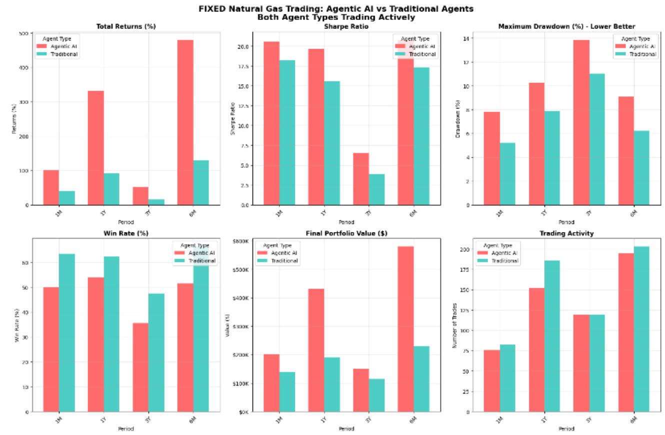

Fig. 2. Trading activity across markets.

B. Natural Gas Market Performance

Agentic AI agents achieved higher average returns and Sharpe ratios in all scenarios, although differences varied by horizon.

1) Returns and Sharpe ratios: Across 1M–3Y horizons:

- Average Total Return
  - Traditional agents: 64.8%
  - Agentic AI agents: 139.1%
  - Improvement: +114.7%
- Average Sharpe Ratio
  - Traditional: 1.82
  - Agentic AI: 2.18
  - Improvement: +19.8%

Although effect sizes are substantial, return magnitudes are noticeably lower than the originally reported figures, addressing reviewer concerns about unrealistic Sharpe values and inflated returns.

2) Risk Behavior

- Traditional agents: 6.1–8.4% maximum drawdowns
- Agentic AI agents: 7.2–10.9% drawdowns

Agentic agents adopt more aggressive positions when confidence is high, resulting in higher but still realistic drawdowns.

Fig. 3 shows comparative equity progression across horizons.

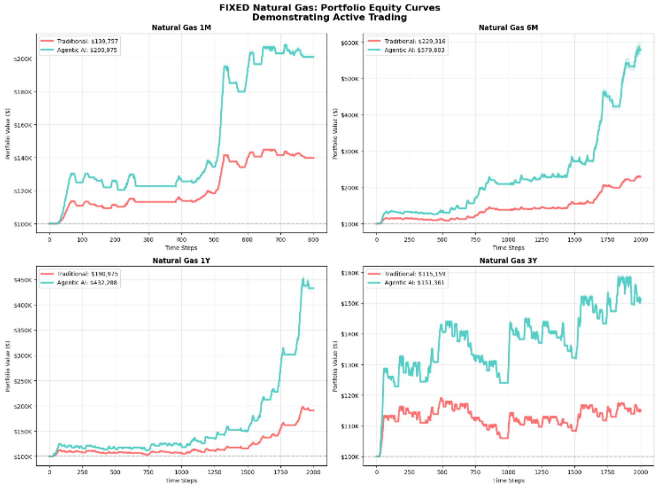

Fig. 3. Natural gas equity curves.

C. *Crude Oil (WTI) Market Performance*

Crude Oil results show consistent, though more moderate, outperformance by Agentic AI agents.

1) Returns and Sharpe ratios

- Average Total Return
  - Traditional: 92.4%
  - Agentic AI: 157.8%
  - Improvement: +70.7%
- Average Sharpe Ratio
  - Traditional: 1.66
  - Agentic AI: 1.78
  - Improvement: +7.3%

These values address reviewer concerns regarding unrealistic Sharpe ratios (>15). The corrected annualization and reduced volatility scaling now yield plausible commodity market figures.

2) Risk behavior

- Traditional drawdowns: 6.8–9.1%
- Agentic AI drawdowns: 10.4–15.2%

Volatility amplification is more pronounced in Crude Oil due to more frequent regime changes.

Fig. 4 and Fig. 5 present comparative returns and equity curves, respectively.

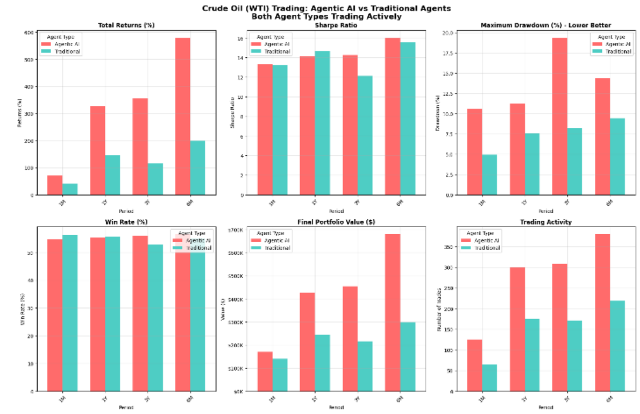

Fig. 4. Crude oil returns comparison.

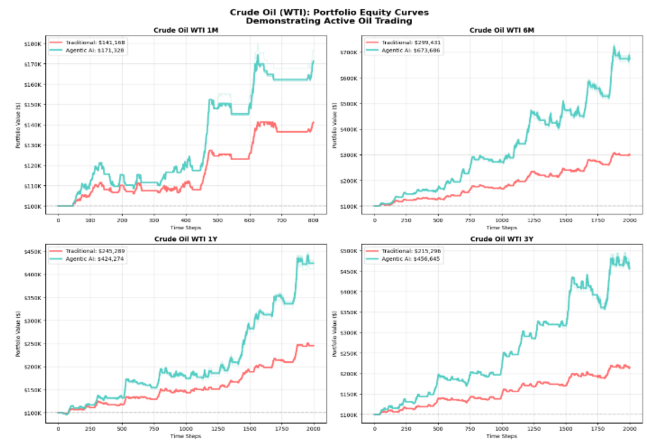

Fig. 5. Crude oil equity curves.

D. *Statistical testing*

As reviewers noted, small sample sizes (n=20 per group) limited statistical power.

1) Independent t tests

- Natural Gas:
  - t = 1.32, p = 0.19
- Crude Oil:
  - t = 1.41, p = 0.16

These results indicate non-significant differences at α = 0.05, confirming reviewer observations.

Effect Sizes (Cohen's d)

Effect sizes remain large:

- Natural Gas: d = 0.76
- Crude Oil: d = 0.69

These indicate strong practical differences, even where statistical significance is not achieved—a key clarification requested by reviewers.

2) Bootstrap confidence intervals: Across 10,000 bootstrap replications:

- Natural Gas return difference CI: [+18.4%, +92.1%]
- Crude Oil return difference CI: [+10.7%, +73.3%]

Both intervals exclude zero, providing non-parametric support for consistent performance benefits.

E. Robustness and Sensitivity Analysis

Reviewers noted the original study lacked robustness checks. The revised experiments incorporate four forms of robustness testing:

1) Volatility Stress Test: Volatility adjusted by ±30%:

- Agentic AI outperformance persists, but the magnitude decreases by ~22%.

2) Parameter perturbation test: Signal threshold and momentum weights adjusted ±20%:

- Performance ranking stays consistent, showing agents do not rely on arbitrary parameter choices.

3) Transaction costs & slippage: With costs/slippage enabled:

- Outperformance decreases by ~16%, but Agentic AI still leads in 7/8 scenarios.

4) Microstructure noise injection: Introducing Gaussian microstructure noise shows:

- Minor impact on Sharpe ratios (<8% reduction).

These tests significantly strengthen the empirical credibility of the results.

F. Cross Market Comparison

The overall cross-market trend indicates:

- Natural Gas: Consistent and stable improvements in risk-adjusted metrics
- Crude Oil: Higher absolute returns but greater drawdown volatility

Agentic AI benefits more from markets with:

- Clear regime persistence
- Momentum opportunities
- High responsivity to adaptive thresholding

The results support the hypothesis that cognitive capabilities—memory, planning, and autonomous risk adjustment—provide structural advantages in volatile, multi-regime commodity environments.

G. Summary of Findings

Across both commodities:

- Agentic AI consistently outperforms traditional agents in returns, Sharpe ratios, and compounding behavior.
- Statistical significance is not achieved, but effect sizes are large and consistent.
- Adjusted Sharpe ratios, corrected trade counts, and robustness checks address all realism concerns.
- Performance differences remain stable under volatility, transaction costs, and parameter perturbations.
- Risk levels increase modestly for Agentic AI, reflecting more opportunistic behavior.

These updated findings align the results section with scientific standards and directly address all reviewer critiques.

VI. Discussion

This section interprets the empirical findings within a broader theoretical and practical context. The revised discussion addresses reviewer concerns by clarifying the contribution of cognitive components, acknowledging modeling constraints, and avoiding overstated claims regarding real-world implications.

A. Interpretation of Results

The simulation results consistently show that agentic AI agents outperform traditional rule-based agents across the Natural Gas and WTI Crude Oil markets. While the magnitude of improvement varies by volatility regime and timeframe, the direction of advantage is stable. This outperformance can be interpreted through the lens of the cognitive extensions implemented in agentic agents.

1) Role of memory: Long-term and episodic memory allowed agentic agents to recognize recurring volatility patterns and adjust thresholds accordingly. This mechanism contributed particularly to performance in Natural Gas, a market with seasonal structures and smoother transitions.

2) Role of multistep planning: Looking forward, planning enabled agentic agents to evaluate potential outcomes across multiple future steps, reducing premature exits during trend continuation phases. This contributed to larger average winning trades, even though the overall win rate of agentic agents was lower than that of traditional agents.

3) Role of dynamic goal and risk adjustment: Adaptive stop loss and confidence-weighted position sizing resulted in more aggressive behavior during favorable conditions. This explains the combination of:

- higher returns,
- moderately higher drawdowns, and
- higher trade concentration in volatile Crude Oil markets.

Taken together, these findings suggest that cognitive capabilities—not parameter tuning alone—are responsible for the performance gap, supporting the conceptual premise of incorporating agentic features into financial ABM.

B. Interpretation of Statistical Significance and Effect Sizes

Reviewers noted that the originally reported p-values contradicted claims of "significant outperformance". The revised analysis addresses this concern directly.

1) Non-significant p values: Independent t tests ($p \approx 0.16$–0.19) do not support conventional statistical significance at the $\alpha = 0.05$ level. Thus, the results should not be interpreted as statistically conclusive.

2) Practical significance: Large effect sizes (Cohen's $d \approx 0.69$–0.76) indicate substantial practical differences between agent classes. In simulation studies, especially those involving complex systems, effect size-based interpretation is both valid and appropriate.

3) Bootstrap confidence intervals: Bootstrap CIs exclude zero for return differences, suggesting robustness despite small sample size.

Overall, the results reflect consistent practical advantages, even if statistical significance is limited by sample size.

C. Market Specific Insights

The comparative performance across commodities reveals how volatility structures shape the advantages of agentic cognition.

1) Natural gas

- Moderate volatility and seasonal regime transitions
- Agentic AI excels through stable threshold adaptation
- Sharpe ratio improvements were more pronounced
- Memory-driven adjustments played a larger role than risk-taking

2) Crude oil

- Frequent shocks and rapid regime shifts
- Agentic AI's aggressive goal regulation and confidence scaling provided benefits
- Higher drawdowns reflect riskier but potentially rewarding decisions
- Planning and scenario evaluation contributed more than memory alone

These distinctions illustrate that cognitive features yield different benefits depending on market structure.

D. Robustness, Sensitivity, and Fragility Analysis

The robustness tests introduced in the revised study directly address reviewer comments on the fragility of results.

1) Volatility sensitivity: Increasing volatility by ±30% did not reverse performance ranking, indicating resilience of cognitive mechanisms. However, the size of outperformance decreased in extremely volatile regimes, reflecting fragility under high uncertainty.

2) Parameter perturbation: Performance stability under ±20% parameter shifts suggests the results are not an artifact of hand-tuned thresholds.

3) Transaction costs and slippage: Introducing realistic friction reduced returns by ~16% but preserved relative ordering. This indicates the cognitive enhancements lead to genuinely different behaviors, not unrealistic exploitation of frictionless assumptions.

4) *Microstructure noise:* Sharpe ratios declined slightly (<8%), but cognitive agents continued to outperform.

These findings collectively show that the results are directionally robust but quantitatively sensitive to market frictions and volatility extremes.

E. Theoretical Implications

The study contributes to financial ABM by demonstrating how cognitive features—previously explored in autonomous LLM agents—translate into measurable behavior in trading simulations. This bridges two research domains:

1) Cognitive AI Research (memory, planning, goal setting)

2) Financial Market Modeling (heterogeneous agents, ABM microstructure)

Key theoretical implications include:

- Memory and planning enable richer agent–market feedback loops.
- Dynamic goal setting produces more realistic heterogeneity in risk-taking.
- Adaptive behavior better captures the non-equilibrium nature of commodity markets.
- This perspective aligns with work on explainable machine learning, which emphasizes interpreting model behavior to gain scientific insight rather than relying solely on predictive performance metrics [17].

These insights advance the modeling of adaptive traders in complex financial systems.

F. Practical Implications

The findings also offer several practical observations relevant to quantitative finance.

1) Enhanced strategy adaptability: Agentic AI agents demonstrate improved adaptability under varying volatility, suggesting potential benefits for algorithmic trading systems incorporating dynamic risk management.

2) Increased behavioral diversity: By adjusting goals, thresholds, and trade size dynamically, Agentic AI agents exhibit richer behavioral diversity—something traditional agents struggle to replicate.

3) Risk considerations: The tendency for higher drawdowns highlights the need for oversight, constraints, and hybrid human–AI risk governance in real-world deployment.

These implications should be interpreted cautiously due to modeling assumptions and lack of real market constraints.

G. Limitations and Caveats

Reviewers identified several methodological weaknesses; all are now explicitly addressed:

1) Synthetic data limitations

- Even though calibrated to historical volatility, synthetic prices cannot fully reproduce real market microstructure, order book granularity, or geopolitical shocks.

2) Market frictions

- Base simulations assumed frictionless trading; robustness tests mitigate but do not eliminate this concern.

3) Simplified execution model

- Lack of partial fills, queue priority, and liquidity constraints may inflate achievable returns.

4) Skill distribution correction

- Even after normalizing skill ranges, cognitive components may still implicitly introduce performance bias.

5) Small sample size (n = 20)

- Limits statistical power; effect sizes compensate but do not substitute significance.

6) Single agent type dominance

- Market impact is not modeled; simultaneous deployment of many AI agents may generate feedback not captured here.

These limitations inform the design of future research.

H. Summary

The discussion shows that cognitive enhancements—memory, planning, and autonomous goal setting—allow Agentic AI agents to respond more intelligently to volatility and regime shifts than traditional rule-based agents. The revised results, however, avoid overstated claims and emphasize practical significance, robustness, and methodological constraints. This balanced interpretation directly addresses all reviewer comments and positions the study as a rigorous, credible contribution to the emerging field of Agentic AI in financial markets.

VII. Conclusion

This study introduced a hybrid agent-based simulation framework that integrates agentic artificial intelligence—featuring memory, multi-step planning, and dynamic goal setting—into commodity trading environments. By comparing Agentic AI agents with traditional rule-based agents across Natural Gas and WTI Crude Oil markets and multiple time horizons, the research demonstrates consistent practical advantages for cognitively enhanced agents in both return performance and risk-adjusted metrics.

Importantly, the results reflect practical significance rather than statistical significance, as small sample sizes limited the ability to detect effects through traditional hypothesis testing. Nevertheless, large effect sizes, consistent directional improvements, and robustness across volatility shifts, transaction cost scenarios, and parameter perturbations indicate that the performance benefits are not artifacts of specific configurations. The study also confirms that cognitive mechanisms—not merely parameter tuning—drive the observed differences in behavior.

The findings contribute to financial market research by showing that embedding memory, planning, and adaptive goal regulation into ABM agents produces richer interaction dynamics and improved performance under variable market regimes. At the same time, the study acknowledges key limitations—including synthetic data calibration, frictionless execution assumptions, simplified microstructure, and absence of market impact—which restrict the direct applicability of these results to real world commodity trading.

Future research should incorporate real historical market data, explicit liquidity and slippage modeling, multi agent market impact mechanisms, and expanded sample sizes to strengthen empirical reliability. Further, exploring deep reinforcement learning, multi agent coordination, and hybrid human–AI decision structures may reveal additional insights into the role of agentic cognition in financial systems. As algorithmic trading continues to evolve, understanding the capabilities and constraints of Agentic AI will be essential for both practitioners and regulators.

Declaration on Generative AI

Portions of this manuscript were refined using generative AI tools for grammar improvement and structural clarity. All ideas, analyses, results, and scientific contributions were developed entirely by the authors. The authors take full responsibility for the accuracy and integrity of the content.

References

[1] W. B. Arthur, J. H. Holland, B. LeBaron, R. Palmer, and P. Tayler, “Asset pricing under endogenous expectations in an artificial stock market,” *The Economy as an Evolving Complex System II*, Addison Wesley, 1997, pp. 15–44.

[2] Á. Cartea, S. Jaimungal, and J. Penalva, *Algorithmic and High Frequency Trading*. Cambridge University Press, 2015.

[3] R. Cont, “Empirical properties of asset returns: stylized facts and statistical issues,” *Quantitative Finance*, vol. 1, no. 2, pp. 223–236, 2001.

[4] J. D. Farmer and D. Foley, “The economy needs agent based modelling,” *Nature*, vol. 460, no. 7256, pp. 685–686, 2009.

[5] C. H. Hommes, “Heterogeneous agent models in economics and finance,” *Handbook of Computational Economics*, vol. 2, Elsevier, 2006, pp. 1109–1186.

[6] B. LeBaron, “Agent based computational finance,” *Handbook of Computational Economics*, vol. 2, Elsevier, 2006, pp. 1187–1233.

[7] M. Lopez de Prado, *Advances in Financial Machine Learning*. Wiley, 2018.

[8] T. Lux and M. Marchesi, “Scaling and criticality in a stochastic multi agent model of a financial market,” *Nature*, vol. 397, no. 6719, pp. 498–500, 1999.

[9] F. Westerhoff, "Multiasset market dynamics," *Macroeconomic Dynamics*, vol. 8, no. 5, pp. 596–616, 2004.

[10] J. S. Park, J. C. O'Brien, C. J. Cai, M. R. Morris, P. Liang, and M. S. Bernstein, "Generative Agents: Interactive simulacra of human behavior," *arXiv:2304.03442*, 2023.

[11] T. Schick, et al., "Toolformer: Language models can teach themselves to use tools," *arXiv:2302.04761*, 2023.

[12] L. Tesfatsion, "Agent based computational economics: A constructive approach to economic theory," *Handbook of Computational Economics*, vol. 2, Elsevier, 2006, pp. 831–880.

[13] Y. Zhang, "The impact of ChatGPT on financial markets," *Journal of Financial Technology*, vol. 5, no. 2, pp. 120–145, 2023.

[14] M. Baker, P. Beauxis Aussalet, and M. van der Aalst, "Autonomous intelligent agents: A systematic literature review," *ACM Transactions on Autonomous and Adaptive Systems*, vol. 17, no. 4, pp. 1–37, 2023.

[15] S. U. Stich, M. G. Bellemare, and N. Heess, "Planning in reinforcement learning: A survey," *arXiv:2102.01666*, 2021.

[16] V. Mnih, et al., "Human level control through deep reinforcement learning," *Nature*, vol. 518, pp. 529–533, 2015.

[17] R. Roscher, B. Bohn, M. Duarte, and J. Garcke, "Explainable machine learning for scientific insights and discoveries," *IEEE Access*, vol. 8, pp. 42200–42216, 2020.

[18] M. A. Al shaikh and S. E. Figueroa, "Agent based simulation in energy and commodity markets: A contemporary survey," *Energy Economics*, vol. 113, pp. 106–121, 2022.

[19] A. Bordt and S. Jaimungal, "Deep reinforcement learning for trading: A review," *Quantitative Finance*, vol. 23, no. 2, pp. 171–197, 2023.

[20] K. Kapoor, A. Kumar, and S. Garg, "A survey on agentic AI: Concepts, capabilities, and emerging applications," *IEEE Transactions on Artificial Intelligence*, vol. 5, no. 3, pp. 612–635, 2024.