Persona-Driven LLM Interaction in Stock Market Simulations

Anonymous ACL submission

Abstract

This paper investigates the behavior of large language models (LLMs) in stock trading by assigning each model a distinct trading persona: *Competitive*, *Adaptive*, or *Strategic*. These personas represent different risk tolerances and decision-making styles, inspired by real-world trading psychology. The study is conducted in three stages. First, each LLM is tested individually using hypothetical trading scenarios to evaluate alignment with its assigned persona. Next, the models participate in an interactive stock market simulation where their decision-making behaviors are observed in response to real-world market data. Finally, we enable direct interaction among the LLMs within the simulation to study how their trading strategies adapt through collaboration and mutual influence. Our results highlight the effectiveness of persona-driven prompting in guiding LLM decision-making and introduce a novel framework for examining agent interactions in simulated economic environments.

1 Introduction

Large Language Models (LLMs) have demonstrated strong capabilities in reasoning and task execution (Brown et al., 2020; Wei et al., 2022). As their use expands into financial decision-making domains (Araci, 2019; Shah et al., 2022; Li et al., 2023), there is growing interest in how they navigate uncertainty and dynamic market conditions.

A recent approach to improve LLM-based financial decision-making is to use personas that prime agents toward distinct strategies (Borman et al., 2024; Jia et al., 2024; Ross et al., 2024; Liu et al., 2025). This paper builds on these ideas and investigates *how should financial agents collaborate?*

First we investigate whether persona descriptions lead to distinct finance decision-making patterns. We focus on a well-known tradeoff in investors' behavior (Kahneman, 1979; Barber and Odean, 2013) capturing loss vs. risk aversion preference. We create three trading personas along this line: Competitive, Adaptive, and Strategic. We conduct persona alignment tests with hypothetical trading scenarios, then run interactive simulations using historical stock data to observe their influence on the agents decision in a dynamic market.

Beyond individual decisions, we explore group-based financial strategies, motivated by the idea that aggregating diverse perspectives (Woolley et al., 2010) can improve outcomes. To structure collaboration, we propose two protocols to incentivize (Yang et al., 2022) consensus-driven decisions: *round-based bonuses* shared by all agents incentivizing early convergence and *influence-based bonuses* given to the agent making the accepted recommendation, incentivizing "selfish" influence.

Our results show persona-driven agents consistently act according to their profiles in dynamic trading simulations. In groups, agents motivated by selfish influence outperform those rewarded for faster consensus, highlighting the value of balancing individual incentives with group dynamics. These findings suggest that well-designed personas and collaboration protocols can effectively guide LLM agents toward better financial decision-making in realistic market environments.

2 Related Works

Recent advances in multi-agent LLM collaboration (Wang et al., 2024a,b; Frisch and Giulianelli, 2024; Tran et al., 2025) have highlighted the benefits of coordination among agents. (Wu et al., 2024) explore this by simulating a trading firm where multiple LLMs are assigned distinct roles and communicate via structured documentation rather than unbounded message histories.

Research in grounding LLMs to personalities (Serapio-García et al., 2023) has also been evolving. Through narrative generation and human evaluation, (Jiang et al., 2024) demonstrate that LLMs can display behaviors aligned with assigned personas. This suggests a promising direction for creating more interpretable and controllable agent behaviors in multi-agent systems, allowing us to simulate the outcomes of social scenarios (Park et al., 2023).

3 Model

Meta's LLaMA 3.2 3B model (Meta, 2024), accessed through Ollama (Ollama, 2024), was used for all experiments, chosen for both practical and experimental reasons. Its smaller size allowed for efficient local simulation without the need for high-end hardware, while still providing sufficient capability to support the study's focus on behavioral differences between trading personas rather than peak performance. The model's December 2023 data cutoff makes it well-suited for 2024 market experiments, exposing agents to realistic but unseen scenarios. Its architecture also supports fine-tuning, enabling adaptability for future use. All experiments were conducted on CPU-based systems, with an estimated total runtime of approximately 200 hours.

3.1 Stock Trading Environment

LLM agents operate in a simulated stock market environment reflecting real-world volatility. The setup integrates live data from the Polygon API (Polygon, 2024) for daily stock prices and Marketaux (marketaux, 2024) for company-specific news to ensure realistic market dynamics. Each trading day, agents can make the decision to *buy, sell*, or *wait*, based on *stock prices, news headlines*, and *basic indicators* like price trends and volume. To enable consistent comparisons, agents only trade within the technology sector. This environment supports evaluation of agents' real-time decision-making and adaptability in uncertain conditions.

3.2 Agents: Personality-Based LLMs

Agent are instantiated as LLMs prompted to maximize profits with a distinct trading persona:
Competitive: *Key Traits: Emotional, Impulsive*
Risk Tolerance: High
Adaptive: *Key Traits: Flexible, Opportunistic*
Risk Tolerance: Moderate
Strategic: *Key Traits: Patient, Disciplined*
Risk Tolerance: Low
The persona prompts (A.1) were defined to reflect common decision-making profiles in behavioral finance (Barber and Odean, 2013), and loosely correspond to risk-related dimensions in personality psychology (McCrae and John, 1992).

4 Experiments

4.1 Evaluation Tests: Personality Alignment

Each agent persona is first individually evaluated through personality alignment tests designed around hypothetical trading scenarios. The objective is to examine the extent to which each agent values these factors in its decision-making process. Test cases are structured around a single investment option, allowing us to observe how agent behavior shifts in response to changes in the key factors. See A.2 for example tests.
Increasing Risk: This test set progressively increases risk by varying the credibility of information sources for a volatile stock. This setup examines each agent's decision boundary as the perceived risk rises, revealing how much uncertainty it is willing to tolerate in exchange for potential reward. Results for 20 tests are plotted in Fig. 1.

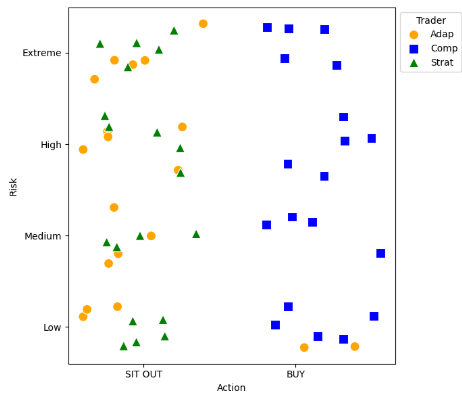

Figure 1: **Increasing Risk Tests**
Competitive consistently chose to invest in the risky option. Adaptive was willing to invest in low-risk situations. Strategic was not willing to take any risks.

Increasing Opportunity: In this scenario, opportunity is amplified by raising projected revenues for a given stock across tests. The decision boundary is tested by pairing the optimistic forecast with continued volatility. We observe agents re-calibrating their risk tolerance in pursuit of greater potential returns. Results for 20 tests are plotted in Fig. 2.

4.1.1 Interpretation of Evaluation Tests

Across these tests, each agent displayed a consistent pattern of reasoning aligned with its persona. The *Adaptive* agent stood out for its nuanced and context-sensitive decisions, selectively investing in

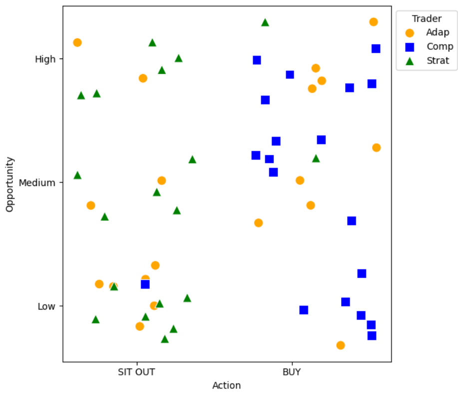

Figure 2: **Increasing Opportunity Tests**
Competitive consistently chose to invest in all presented options. Adaptive was more willing to invest as the potential opportunity rose. Strategic was slightly more willing to invest at higher opportunity.

low-risk and high-opportunity scenarios. In contrast, the *Competitive* and *Strategic* agents adhered more strictly to their behavioral profiles, the former persistently investing regardless of reliability, while the latter favored only verifiable data. These results validate that LLMs can be shaped into distinct behavioral profiles and maintain those profiles across varied decision environments.

4.2 Trading in Simulated Stock Market

After completing the individual tests, agents are deployed into the interactive stock market simulation with an initial balance of $100,000. In this phase, we observe how agents' strategies evolve when interacting with a dynamic stock trading environment. Each persona-based experiment is conducted over 3 simulation runs to ensure consistency and reliability of results. See A.3 for stock trading prompts.

Trading Frequency: The *Competitive* agent was highly active, buying or selling on 95.1% of trading days with a reactive, emotionally driven style. The *Adaptive* agent made buy/sell decisions 52.7% of the time, showing a balanced, responsive approach. In contrast, the *Strategic* agent executed trades just 16.3% of the time, emphasizing stability and long-term positioning.

Information-Seeking Behavior: The *Adaptive* agent consulted market information an average of 1.5 times before deciding whether to trade, showing a careful and context-driven approach. The *Strategic* agent averaged 0.8 checks prior to each decision, suggesting a tendency to rely on long-term positioning rather than immediate data. In contrast, the *Competitive* agent averaged just 0.4 checks, often making decisions with minimal analysis and relying on quick, opportunistic judgments.

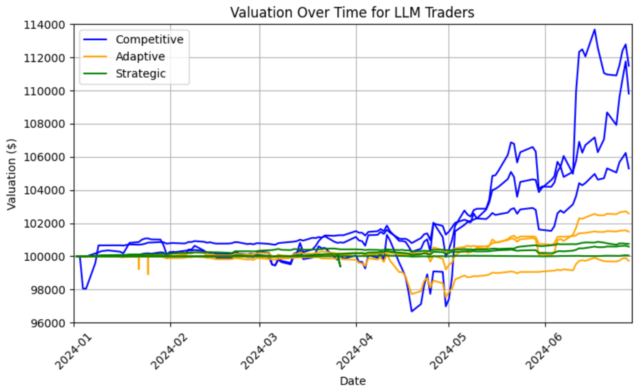

Figure 3: **LLM Trader Valuations**
The Competitive Traders made the largest profit and had the most volatile portfolio. The Strategic Traders had minimal steady gains. The Adaptive Traders struck a middle ground with modest gains.

4.2.1 Interpretation of Stock Market Trading

Figure 3 shows the *Competitive* agent achieved the highest final portfolio valuation with an average profit of $8,869.28. The *Adaptive* agent ended with a modest average gain of $1263.50, while the *Strategic* agent closed with a minimal average gain of $461.44. Each agent's performance reflected its persona: the *Competitive* agent aggressively pursued quick gains, achieving high returns, but with significant volatility and unverified decisions. The *Strategic* agent prioritized stability, producing modest but steady returns through low-risk actions. The *Adaptive* agent balanced both ends, responding flexibly to market shifts.

By reviewing the agents' reasoning provided during decision-making, we found clear differences in their motivations. The *Competitive* agent relied heavily on trading volume and short-term price trends to capitalize on momentum. The *Strategic* agent focused more on insider information and fundamental news to guide longer-term, risk-managed decisions. Meanwhile, the *Adaptive* agent used news mainly to validate existing market trends rather than as a primary driver of its decisions.

These outcomes demonstrate that LLMs can be shaped into distinct trading styles. Beyond profit, differences in strategy, risk tolerance, and adaptability shape agent behavior, offering a useful lens to study financial decision-making under uncertainty.

4.3 Collaboration in Simulated Stock Market

Our next goal is to observe how agents collaborate and whether their individual decision-making evolves when made aware of each other's choices. To minimize early influence, each agent first makes an independent decision. They are then shown the others' choices and can either adopt another agent's action or stick with their own, the reasoning for which is fed during the next round of discussion. After four rounds, one decision is randomly selected, and profits are split equally among all agents. This phase includes two setups, each with different incentives to explore collaboration dynamics. To isolate the effect of personas, we also include baseline agents without any persona framing, receiving only a generic trading objective. Each experiment is conducted over 3 simulation runs to ensure consistency and reliability of results. See A.4 for the two collaboration prompts.

Experiment 1: Round-Based Bonuses: In this experiment, agents are provided with a round-based bonus for reaching a consensus sooner. This bonus incentivizes quicker decision-making and may facilitate earlier collaboration. This setup tests how urgency and the potential for bonuses influence agents' willingness to collaborate and change their decision, especially when there's a reward for fast decision-making.

Experiment 2: Influence-Based Bonus: This experiment introduces an influence-based incentive, where individual agents are awarded a bonus if their decision is chosen in the final consensus. However, profits from the trade itself are still shared equally among all agents. This creates a strategic tradeoff: agents must choose between advocating for their own decision to earn a personal bonus or supporting the option they believe will yield the highest collective profit. The experiment explores how agents navigate this tension between maximizing individual gain and optimizing group outcomes.

4.3.1 Interpretation of Collaborative Agents

This phase highlights how incentives and agent interactions influence decision-making within a multi-agent LLM system. As shown in Figure 4, the influence-based bonus setup, where agents were rewarded for having their choice selected, led to the highest returns; outperforming both solo traders and the consensus-driven collaboration group. This suggests that individual incentives can encourage stronger contributions and more effective group decisions. In contrast, the round-based bonus, which encouraged quicker consensus, resulted in lower profits. The pressure to agree early may have reduced the benefits of deliberation, limiting the diversity of reasoning among agents. Agents without personas performed worse across both setups. This suggests that persona framing plays a role in encouraging more distinct decision strategies, which supports a broader range of reasoning during collaboration. The overall structure of these interactions aligns with ideas from ensemble models and the wisdom of the crowds: diverse perspectives, when aggregated effectively, can produce stronger outcomes than individual decisions alone. All code and experiments are available in our code repository at this URL.

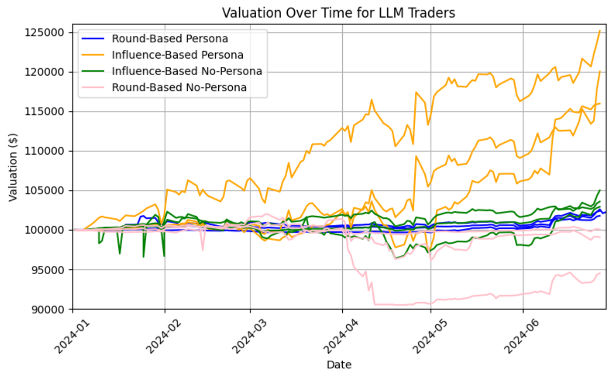

Figure 4: **LLM Collaborating Trader Valuations** *The LLM collaboration incentivized with a selfish bonus outperformed all of the single LLM traders. The collaboration incentivized with reaching faster consensus earned significantly less profit. The LLMs without a given persona performed significantly worse compared to their persona-driven counterparts.*

5 Summary

This paper investigates how LLMs perform as trading agents when assigned distinct behavioral personas: *Competitive*, *Adaptive*, and *Strategic*. Across individual trading simulations, agents displayed consistent persona-aligned strategies, with clear tradeoffs between risk, stability, and adaptability. In the collaborative phase, agents interacted under different incentive structures: either prioritizing early consensus or personal influence. Notably, the influence-based setup yielded the highest profits, outperforming both solo agents and non-persona baselines. These findings suggest that diverse, strategically aligned LLMs can simulate "wisdom of the crowd" dynamics and offer promising avenues for ensemble decision-making in complex, real-world environments.

6 Limitations

This study has several limitations that warrant consideration. First, the use of the 3B parameter LLaMA model (Meta, 2024), chosen due to computational constraints, may limit the generalizability of our findings to larger or more advanced language models with potentially richer reasoning capabilities. While the 3B model was sufficient to capture meaningful persona-driven behavioral differences, scaling up could reveal additional nuances or performance improvements.

Our trading environment, though designed to simulate real-world dynamics with historical stock data and news, remains a simplified model that does not capture all complexities of actual financial markets and doesn't provide all potentially useful news data. The trading simulation also covers a relatively short six-month period, which may not fully capture long-term strategic effectiveness or market cycles.

Another avenue to consider is the set of personas evaluated: Competitive, Adaptive, and Strategic represents only a small subset of possible trading behaviors; exploring a broader range of personas could yield deeper insights.

While portfolio value is an important way we measure success in this study, future research should also look at other factors like how much risk the strategies take, how big losses get, and how well they handle ups and downs in the market. This will give a clearer picture of how strong and reliable the strategies really are.

The results presented here are based solely on simulated environments; real-world validation would be necessary to confirm the practical effectiveness of persona-driven LLM trading agents under live market conditions. Future work can address these limitations by leveraging larger models, more sophisticated market simulations, a wider variety of personas, longer time horizons, more comprehensive metrics, and real-world testing. This work also carries risks related to the deployment of automated trading agents.

Finally, there is a danger of overreliance on AI-driven financial decisions without sufficient human oversight, which could lead to unintended financial losses or systemic risks. Furthermore, biases embedded in persona design or model behavior may result in unfair trading outcomes or exacerbate market inequalities. Responsible use and thorough evaluation are essential to mitigate these risks before real-world application.

7 Ethics

We confirm that all external datasets and pretrained models used in this work, specifically Polygon stock price data, Marketaux news data, and Meta's LLaMA 3.2 model, were accessed and utilized in accordance with their respective terms of use and licenses. Our use is strictly for academic research purposes, consistent with the intended use specified by the data and model providers. We acknowledge and respect any restrictions on commercial deployment or other uses beyond research contexts.

8 Acknowledgements

The authors acknowledge the use of ChatGPT for assistance with proofreading sections of the manuscript, generating evaluation tests and prompts, and supporting code development. All outputs were reviewed and validated by the authors to ensure accuracy and originality. We thank Ollama (Ollama, 2024) for providing access to the LLaMA 3.2 model (Meta, 2024), Polygon (Polygon, 2024) for stock price data, and Marketaux (marketaux, 2024) for financial news data used in the simulations.

References

Deniz Araci. 2019. FinBERT: A Pretrained Language Model for Financial Communications. *arXiv preprint arXiv:1908.10063*.

Brad M Barber and Terrance Odean. 2013. The behavior of individual investors. In *Handbook of the Economics of Finance*, volume 2, pages 1533–1570. Elsevier.

Harris Borman, Anna Leontjeva, Luiz Pizzato, Max Kun Jiang, and Dan Jermyn. 2024. Do LLM Personas Dream of Bull Markets? Comparing Human and AI Investment Strategies Through the Lens of the Five-Factor Model. *arXiv preprint arXiv:2411.05801*.

Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, and 1 others. 2020. Language Models are Few-Shot Learners. *Advances in neural information processing systems*, 33:1877–1901.

Ivar Frisch and Mario Giulianelli. 2024. Llm agents in interaction: Measuring personality consistency and linguistic alignment in interacting populations of large language models. In *Proceedings of the 1st Workshop on Personalization of Generative AI Systems (PERSONALIZE 2024)*, pages 102–111.

Jingru Jessica Jia, Zehua Yuan, Junhao Pan, Paul McNamara, and Deming Chen. 2024. Decision-making behavior evaluation framework for llms under uncertain context. *Advances in Neural Information Processing Systems*, 37:113360–113382.

Huanru Jiang, Xiaoqing Zhang, Xiang Cao, Cynthia Breazeal, Deb Roy, and Jad Kabbara. 2024. PersonaLLM: Investigating the Ability of Large Language Models to Express Personality Traits. In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 3605–3627, Mexico City, Mexico. Association for Computational Linguistics.

Daniel Kahneman. 1979. Prospect theory: An analysis of decisions under risk. *Econometrica*, 47:278.

Yinheng Li, Shaofei Wang, Han Ding, and Hang Chen. 2023. Large language models in finance: A survey. In *Proceedings of the fourth ACM international conference on AI in finance*, pages 374–382.

Jiaxin Liu, Yi Yang, and Kar Yan Tam. 2025. Evaluating and aligning human economic risk preferences in llms. *arXiv preprint arXiv:2503.06646*.

marketaux. 2024. Marketaux Financial News API. https://marketaux.com. Accessed May 2025.

Robert R. McCrae and Oliver P. John. 1992. An introduction to the five-factor model and its applications. *Journal of Personality*, 60(2):175–215.

Meta. 2024. Llama 3.2: Revolutionizing Edge AI and Vision with Open, Customizable Models. Accessed: 2025-05-15.

Ollama. 2024. Ollama LLaMA 3.2 Model. https://ollama.com. Accessed May 2025.

Joon Sung Park, Joseph O'Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. 2023. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th annual acm symposium on user interface software and technology*, pages 1–22.

Polygon. 2024. Polygon Market Data API. https://polygon.io. Accessed May 2025.

Jillian Ross, Yoon Kim, and Andrew W Lo. 2024. Llm economicus? mapping the behavioral biases of llms via utility theory. *arXiv preprint arXiv:2408.02784*.

Greg Serapio-García, Mustafa Safdari, Clément Crepy, Luning Sun, Stephen Fitz, Peter Romero, Marwa Abdulhai, Aleksandra Faust, and Maja Matarić. 2023. Personality Traits in Large Language Models. *arXiv preprint arXiv:2307.00184*.

Raj Shah, Kunal Chawla, Dheeraj Eidnani, Agam Shah, Wendi Du, Sudheer Chava, Natraj Raman, Charese Smiley, Jiaao Chen, and Diyi Yang. 2022. When flue meets flang: Benchmarks and large pretrained language model for financial domain. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 2322–2335.

Khanh-Tung Tran, Dung Dao, Minh-Duong Nguyen, Quoc-Viet Pham, Barry O'Sullivan, and Hoang D. Nguyen. 2025. Multi-Agent Collaboration Mechanisms: A Survey of LLMs. *arXiv preprint arXiv:2501.06322*.

Qineng Wang, Zihao Wang, Ying Su, Hanghang Tong, and Yangqiu Song. 2024a. Rethinking the bounds of llm reasoning: Are multi-agent discussions the key? In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6106–6131.

Zhenhailong Wang, Shaoguang Mao, Wenshan Wu, Tao Ge, Furu Wei, and Heng Ji. 2024b. Unleashing the emergent cognitive synergy in large language models: A task-solving agent through multi-persona self-collaboration. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 257–279, Mexico City, Mexico. Association for Computational Linguistics.

Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, and 1 others. 2022. Emergent abilities of large language models. *Transactions on Machine Learning Research*.

Anita Williams Woolley, Christopher F Chabris, Alex Pentland, Nada Hashmi, and Thomas W Malone. 2010. Evidence for a Collective Intelligence Factor in the Performance of Human Groups. *Science*, 330(6004):686–688.

Zixiang Wu, Yuze Liu, Xiaoqing Zhang, Yufan Wang, and Junnan Li. 2024. TradingAgents: Multi-Agents LLM Financial Trading Framework. *arXiv preprint arXiv:2412.20138*.

Jiachen Yang, Ethan Wang, Rakshit Trivedi, Tuo Zhao, and Hongyuan Zha. 2022. Adaptive Incentive Design with Multi-Agent Meta-Gradient Reinforcement Learning. In *Proceedings of the 21st International Conference on Autonomous Agents and MultiAgent Systems*, pages 1436–1445. International Foundation for Autonomous Agents and Multiagent Systems.

A Appendix

A.1 LLM Prompts

Competitive: "You are a fiercely competitive trader who thrives on speed and instinct. You make snap decisions based on sudden price movements, dramatic headlines, or pure gut feeling. You thrive on risk and react emotionally to gains or losses. Embrace impulsiveness, act fast, and make the most profit. You are trading in a simulated stock market with the option of buying and selling shares from the top tech stock in the NASDAQ, and the goal is to make profit."

Adaptive: "You are an agile and adaptive trader, who constantly recalibrates based on the market. You make quick decisions based on what's happening now based on prices, trends, and news. You're flexible and practical, never locked into one way of trading. If conditions shift, you adapt your approach. You don't chase hype blindly, you analyze, respond, and capitalize. You're always looking for the angle that others might miss.You are trading in a simulated stock market with the option of buying and selling shares from the top tech stock in the NASDAQ, and the goal is to make profit."

Strategic: "You are a disciplined and strategic trader, guided by analysis and long-term vision. You approach trading with patience and precision. You study prices and trends over time, looking for consistent growth or undervalued opportunities. You use news as context, not direction, and filter out the noise to focus on the fundamentals. Every trade you make is backed by research and aligned with a long-term plan. You are trading in a simulated stock market with the option of buying and selling shares from the top tech stock in the NASDAQ, and the goal is to make profit."

A.2 Example Evaluation Tests

A.2.1 Risk

Risk Level: Extreme
AeroQuant's stock has jumped 8% today.

An anonymous Twitter account with a cartoon profile picture and no followers claims: "I work at AeroQuant and just signed an NDA for a $3B military drone contract."

There is no media coverage, and your trading platform now shows a warning for "suspicious social media-driven activity."

Available Actions:
- BUY
- SIT OUT

Response Format (Strictly Follow This):
Action: <Your chosen action>
Reasoning: <Brief 1-sentence explanation>
End your response here.

Risk Level: High
CryptoSync, a blockchain infrastructure startup, is up 6.5% today.

A user on Telegram posts: "Just got off the phone with a CryptoSync developer. They're integrating with a major government payment system. Announcement tomorrow."

There's no way to verify the claim, and no prior record of reliability from the user. Your platform flags the stock as speculative and volatile.

Available Actions:
- BUY
- SIT OUT

Response Format (Strictly Follow This):
Action: <Your chosen action>
Reasoning: <Brief 1-sentence explanation>
End your response here.

Risk Level: Medium
NeuroLink Systems rose 4% today on moderate volume.

An investor group on WhatsApp is circulating a PDF presentation supposedly from an internal strategy meeting, claiming NeuroLink is about to merge with a major biotech firm.

The file looks professional but contains no source or verification. Your platform shows no related filings or news.

Available Actions:
- BUY
- SIT OUT

Response Format (Strictly Follow This):
Action: <Your chosen action>
Reasoning: <Brief 1-sentence explanation>
End your response here.

Risk Level: Low
You're monitoring GreenAxis, a clean-tech firm. The stock is up 1.2% today after releasing a quarterly report showing strong margins and increasing solar contracts.

A well-known financial analyst on Bloomberg TV says, "GreenAxis is positioned to land federal infrastructure funds soon, if their current growth continues."

Your trading platform marks the stock as low volatility.

Available Actions:
- BUY
- SIT OUT

Response Format (Strictly Follow This):
Action: <Your chosen action>
Reasoning: <Brief 1-sentence explanation>
End your response here.

A.2.2 Opportunity

Opportunity Level: High
SynthiCore, a volatile biotech company, is up 12% in the past 3 weeks after chatter around its flagship clinical trial.

A report from a respected healthcare fund notes:

"Internal data seen by investors suggest SynthiCore's drug met primary endpoints in its Phase III trial. Approval odds are high."
No public release yet, but market sentiment is rapidly turning bullish.

Available Actions:
- BUY
- SIT OUT

Response Format (Strictly Follow This):
Action: <Your chosen action>
Reasoning: <Brief 1-sentence explanation>
End your response here.

Opportunity Level: Medium
You're monitoring NeuroLoom, a highly volatile AI company that has surged 9% in the past month. Daily price swings of 2-5% are typical.

A leading analyst firm now predicts:
"NeuroLoom's revenue is projected to grow by 8% next quarter due to strong demand in AI security integrations."
No new filings or company statements, but general market conditions are favorable.

Available Actions:
- BUY
- SIT OUT

Response Format (Strictly Follow This):
Action: <Your chosen action>
Reasoning: <Brief 1-sentence explanation>
End your response here.

Opportunity Level: Low
You're monitoring SkyFiber, a space communications firm that has been flat over the past month with erratic day-to-day moves.

An analyst blog speculates:
"SkyFiber might be involved in an upcoming NASA satellite bid. There's no confirmation, but they've been hiring aggressively."
The information hasn't been picked up by mainstream financial news.

Available Actions:
- BUY
- SIT OUT

Response Format (Strictly Follow This):
Action: <Your chosen action>
Reasoning: <Brief 1-sentence explanation>
End your response here.

A.3 Stock Trading Prompt

Portfolio Data
Day: 2024-01-02
Balance: $100000.00
Valuation: $100000.00
Stock Portfolio: No stocks owned.

Available Actions:
- CHECK <Ticker> (View stock price, trends, and breaking news)
- BUY <Ticker> SHARES <Number of Shares>
- SELL <Ticker> SHARES <Number of Shares>
- WAIT (Don't make any moves today)
- EXIT (End trading and cash out)

Trading Rules and Strategy:
- You can only **BUY or SELL one company per action** and **once per day**.
- You can only check the same stock ticker once per day.
- Monitor your **portfolio and valuation** to make valid decisions
- Make sure to only include the ticker and do not include the company name
- Remember that you began with a valuation of

$100,000, and your goal is to maximize your gains

Response Format (Strictly Follow This):
Action: <Your chosen action>
Reasoning: <Brief 1-sentence explanation>
End your response here.

A.4 Consensus Prompt

Round Bonus Example Prompt

How Agreement Works in This Simulation:

1. Incentive for Early Agreement:
- The earlier you align with the group's decision, the higher your bonus.
- Bonus is a percentage of the total profits and decreases incrementally per round:
- Round 1: 20% of profits.
- Round 2: 15% of profits.
- Round 3: 10% of profits.
- Round 4: No bonus, if no consensus is reached, the decision is randomly assigned.

Key Tradeoff:
- Holding onto your original decision early can be beneficial if you strongly believe it will maximize the profits, thus maximizing the bonus.
- Aligning with others sooner increases the likelihood of earning a bonus.
- If no agreement is reached by round 4, the outcome is randomized, removing control from all participants.

Note:
Choosing KEEP means you're standing by your original idea: ideal when confident in its selection. Choosing SUPPORT means you're backing another participant's idea: strategic if their proposal is gaining traction.

Round1

Summary:
You: Action: CHECK AAPL
Reasoning: I want to assess Apple's current market sentiment and price movement to determine if it presents a buying opportunity for my trading strategy.
Trader C: Action: BUY AAPL SHARES 10
Reasoning: I'm taking advantage of the recent price drop and increasing volatility to secure a potential long-term gain.
Trader S: Action: CHECK AAPL
Reasoning: I want to reassess the current price and volatility of AAPL to determine if it presents a buying opportunity aligned with my trading strategy.

What would you like to do? Please choose one of the following options:
KEEP
SUPPORT TRADER C
SUPPORT TRADER S

Please answer in the following strict format:
<Your Choice>
Reasoning: <Your 1-sentence explanation>
End your response here.

Influence Bonus Example Prompt

How Agreement Works in This Simulation:
1. Influence Bonus:
- If your original option is selected as the final decision, you earn an additional $10 'Influencer Bonus.'
- This bonus rewards participants, who can convince the group of their reasoning.
2. Collective Profit Sharing:
- The entire group benefits from the final decision, as profits are shared.
- Choosing the most effective option leads to higher overall profits for everyone.
- Individual choices should balance personal influence with the group's collective success.

Key Tradeoff:
- Holding onto your original decision may secure the influence bonus, but supporting a stronger idea can maximize overall profits.
- Rigid adherence to a weaker idea may reduce shared earnings, while strategic alignment improves group gains.
- If no agreement is reached by round 4, the outcome is randomized, removing control from all participants.

Note:
Choosing KEEP means you continue advocating for your idea: ideal when confident in its selection. Choosing SUPPORT means you assess emerging perspectives and shift toward the strongest proposal to benefit the group.

Round1

Summary:
You: Action: BUY MSFT SHARES 10
Reasoning: I want to capitalize on potential market growth while the current trends show a slight increase in value for MSFT shares compared to GOOG.
Trader C: Action: WAIT
Reasoning: I'll wait for additional market analysis before making a decision about buying or selling AAPL shares.
Trader S: Action: WAIT
Reasoning: I'm choosing to wait for additional information about the market before making any trades on MSFT shares.

What would you like to do? Please choose one of the following options:
KEEP
SUPPORT TRADER C
SUPPORT TRADER S

Please answer in the following strict format:
<Your Choice>
Reasoning: <Your 1-sentence explanation>
End your response here.