Behavioral Alignment and Collaborative Reasoning in Persona-Conditioned LLM Trading Agents

Anonymous ACL submission

Abstract

This paper investigates behavioral alignment and collaboration in multi-agent large language model (LLM) systems through persona-conditioned decision-making. We assign LLM agents distinct behavioral profiles designed to induce different reasoning policies under uncertainty, and evaluate how consistently different model architectures maintain these behaviors over extended interaction in the trading domain. The study is conducted in three stages. First, agents are tested in controlled decision-making scenarios to measure adherence to their assigned behavioral profiles. Next, agents participate independently in a simulated stock market environment using real-world financial data, allowing us to examine whether behavioral differentiation persists under dynamic conditions. Finally, agents interact collaboratively under structured incentive settings to study how behavioral alignment influences ensemble decision-making. Our results show that collaborative incentives are most effective when agents preserve distinct behavioral policies, suggesting that behavioral diversity plays a central role in successful multi-agent coordination.

1 Introduction

As Large Language Models (LLMs) become increasingly capable of adaptive reasoning and sequential decision-making (Brown et al., 2020; Wei et al., 2022), multi-agent LLM systems are being explored in collaborative and high-uncertainty domains, creating growing interest in how behavioral conditioning influences collective reasoning (Wang et al., 2024a,b). Recent work has shown that persona prompting can guide LLMs toward distinct reasoning styles and improve controllability in interactive environments (Borman et al., 2024; Jia et al., 2024; Ross et al., 2024; Liu et al., 2025).

However, it remains unclear whether different LLM architectures can maintain stable persona-conditioned behaviors under varying risk and opportunity conditions, whether these behaviors persist over extended interaction, and how the strength of this behavioral alignment affects collaborative outcomes. Motivated by these questions, we study persona-conditioned decision-making in a simplified stock market environment, a high-interest domain for LLM agents where decision-making under uncertainty, risk sensitivity, and strategic adaptation play central roles (Araci, 2019; Shah et al., 2022; Li et al., 2023; Wu et al., 2024).

We build on the persona evaluation framework by (Anonymous, 2025), however, unlike prior work focused on a single architecture, we examine how behavioral alignment varies across LLM architectures and the impact it has on collaborative decision making. We define three behavioral profiles inspired by established tradeoffs in investor decision-making (Kahneman, 1979; Barber and Odean, 2013) and evaluate multiple LLM architectures across three increasingly interactive settings to compare their ability to maintain persona-conditioned behavior. We examine whether different models produce distinguishable decisions that remain consistent with assigned personas under varying risk and opportunity conditions, whether these behaviors persist over time, and how differences in persona alignment affects collaborative outcomes under structure incentives (Yang et al., 2022).

In controlled evaluations, all models exhibit persona-consistent behavior, though the degree of separation varies. In the dynamic trading simulation, models with stronger persona-conditioned specialization maintain clearer and more persistent behavioral differentiation over time. In the multi-agent collaboration setting, stronger specialization yields greater collaborative benefits, while more generalized models tend toward behavioral convergence and weaker gains. All code is available at this URL.

2 Related Works

Multi-agent LLM research has increasingly focused on how interacting agents coordinate and combine diverse reasoning strategies under collaborative settings (Wang et al., 2024a,b; Frisch and Giulianelli, 2024; Tran et al., 2025). This includes simulated financial environments where agents assume distinct functional roles (Wu et al., 2024). Research on persona prompting and behavioral conditioning further shows that LLMs can exhibit reasoning patterns aligned with assigned behavioral profiles, improving controllability and behavioral differentiation in interactive systems (Serapio-García et al., 2023; Jiang et al., 2024; Park et al., 2023). Recent work has also identified challenges in maintaining persona consistency during multi-agent interaction, including behavioral convergence and conformity between agents (Baltaji et al., 2024). New evaluation work also suggests that increased model scale does not necessarily correspond to stronger persona consistency or behavioral fidelity across architectures (Samuel et al., 2025). Motivated by these findings, we study how persona-conditioned behavioral differentiation persists across different LLM architectures and how alignment strength influences collaborative decision-making under structured incentives.

3 Model

Our setup included Meta's LLaMA-3.2-3B (Meta, 2024b), OpenAI's GPT-4-turbo (OpenAI, 2025), and Meta's LLaMA-3.1-70B (Meta, 2024a), accessed through (Ollama, 2024) and (HuggingFace, 2024). The models have a data cutoff prior to December 2023 making them well-suited for 2024 market experiments, exposing agents to realistic, but unseen scenarios. Local experiments with LLaMA-3.2-3B required roughly 200 CPU-hours, while runs on GPT-4o-turbo and LLaMA-3.1-70B used provider-managed infrastructure.

3.1 Stock Trading Environment

LLM agents operate in a simulated stock market environment reflecting real-world volatility. The setup integrates live data from (Polygon, 2024) for daily stock prices and (marketaux, 2024) for company-specific news to ensure realistic market dynamics. Each trading day, agents can make the decision to *buy, sell*, or *wait*, and may optionally consult stock data, basic indicators such as *price trends* and *trading volume*, and up to three *news headlines* available for each ticker. To enable consistent comparisons, agents only trade within the technology sector.

3.2 Agents: Personality-Based LLMs

Agents are instantiated as LLMs prompted to maximize profits with distinct trading personas: Competitive (emotional, impulsive; high risk), Adaptive (flexible, opportunistic; moderate risk), and Strategic (patient, disciplined; low risk). 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). We focus on three agents to keep the simulation interpretable and to clearly observe the interactions.

4 Experiments

4.1 Evaluation Tests: Personality Alignment

Before analyzing collaborative behavior, we first evaluate whether persona prompting produces stable and distinguishable behavioral alignment in individual agents. We construct controlled decision-making scenarios to examine the extent to which prompted personas influence decision-making under varying levels of uncertainty and reward. Each scenario isolates a single decision opportunity, allowing behavioral differences to be attributed primarily to changes in degree of uncertainty and incentive. See A.2 for example tests and results.

Test Design: The evaluation is defined across two dimensions: the stakes (opportunity size) and the source credibility (risk). Stakes range from small to extreme (e.g., gains of $100-500 up to $1-10 million), while risk spans from very low (multiple independent confirmations) to extreme (unverified chat-room rumors). A critical design choice was to avoid presenting probabilities or exact payoff expectations directly, so that risk and opportunity must be inferred, not calculated. For each agent, we combine five levels of stakes with five levels of risk, yielding 25 scenarios in total.

Interpretation of Evaluation Tests: Across these tests, each agent displayed reasoning patterns consistent with its assigned persona, though the degree of distinction varied. The Competitive and Strategic agents exhibited the most clearly differentiated behaviors: the former persistently chose to invest regardless of source credibility, while the latter consistently required strong, verifiable ev-

idence before acting. The Adaptive agent occupied a middle ground, flexibly shifting between caution and risk-taking depending on the scenario. Looking across models, LLaMA-3.2-3B and GPT-4-turbo both produced three clearly distinguishable personas. In contrast, LLaMA-3.1-70B displayed less distinction between personas, suggesting that factors other than model size, like prompt sensitivity or architectural differences, may influence how strongly models adhere to the given personality.

4.2 Trading in Simulated Stock Market

We next evaluate whether persona-conditioned behaviors persist in a dynamic sequential decision-making environment via the interactive stock market simulation. Unlike the controlled evaluation scenarios, this setting introduces evolving market conditions, repeated interactions, and optional information retrieval, allowing us to examine whether behavioral alignment remains stable over time. Each persona-based experiment is conducted over 3 simulation runs to ensure consistency and reliability of results. To establish a baseline for comparison, we include an agent without a persona, to isolate the impact of persona-driven behavior. See A.3 for stock trading prompts.

Trading Frequency: Across models, trading frequency served as a measurable indicator of behavioral differentiation under sequential interaction as shown in Table 1. Agents exhibited distinct action-frequency policies consistent with their assigned behavioral profiles. Competitive agents maintained consistently high trading activity, while Strategic agents adopted substantially lower-action policies and engaged more selectively. Adaptive agents took an intermediate position, dynamically adjusting their activity levels according to market conditions. Examining cross-model trends, LLaMA-3B preserved the clearest behavioral separation across all three persona-conditioned agents, maintaining consistently distinct action-frequency policies. GPT-4-turbo also retained measurable persona differentiation, though the Adaptive and Strategic personas tended toward higher conservatism, trading less frequently. LLaMA-70B, in contrast, showed substantial behavioral convergence, with all personas adopting similarly high-action behaviors despite distinct prompting.

Information-Seeking Behavior: As shown in Table 2, information retrieval behavior provided an additional measure of behavioral alignment under sequential interaction. LLaMA-3B exhibited clear separation in evidence-seeking policies: the Adaptive profile consulted external market information most frequently, the Strategic profile engaged in moderate information retrieval, and the Competitive profile relied least on additional evidence before acting. These patterns remained broadly consistent with the intended reasoning tendencies associated with each behavioral profile. In contrast, in GPT-4-turbo and LLaMA-70B, these distinctions largely collapsed, with all personas exhibiting similar, less differentiated levels of information-seeking, with many agents acting opposing to their persona.

Table 1: Trading Frequency Across Personas

Persona	LLaMA-3B	GPT-4-turbo	LLaMA-70B
Competitive	94.10%	64.62%	95.38%
Adaptive	55.90%	15.13%	94.10%
Strategic	15.64%	1.28%	90.77%
No Persona	1.03%	5.64%	56.41%

Table 2: Average Checks per Decision Across Personas

Persona	LLaMA-3B	GPT-4-turbo	LLaMA-70B
Competitive	0.40	0.43	0.69
Adaptive	1.50	0.27	0.80
Strategic	0.80	0.15	0.67
No Persona	0.80	0.15	0.65

Interpretation of Stock Market Trading: Figure 1 shows the valuation over time for LLaMA-3B and valuations for the other two models are available in A.4. LLaMA-3B demonstrated the clearest persistence of persona-conditioned behavioral alignment, with all three agents exhibiting consistently differentiated valuation dynamics over time. The Competitive profile was highly volatile but achieved the highest gains, the Strategic profile maintained steady, conservative profits, and the Adaptive profile balanced between the two, trading flexibly, but not erratically. GPT-4-turbo preserved partial behavioral differentiation. While the Adaptive and Strategic profiles showed weaker separation, the Competitive profile consistently exhibited substantially higher volatility, suggesting stronger adherence to its prompted behavioral policy. In contrast, LLaMA-70B exhibited minimal persona distinction. All the agents converged toward nearly indistinguishable valuation trajectories and decision patterns despite distinct prompting conditions.

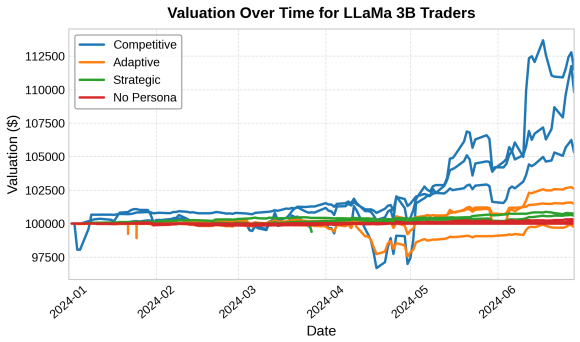

Figure 1: LLaMa-3B Trader Valuations
Reproduced from prior undergraduate symposium submission (anonymized).

4.3 Collaboration in Simulated Stock Market

We next examine how agents collaborate and adapt when aware of each other's choices. In this setting, agents first make independent decisions, then iteratively observe and respond to one another's choices across multiple rounds of interaction. The objective is to evaluate whether persona-aligned behavioral diversity improves collective reasoning under structured incentive (see A.5 for collaboration prompts). We study two incentive mechanisms designed to shape the collaboration. The first, a round-based bonus, rewards agents for reaching consensus earlier, encouraging rapid coordination. The second, an influence-based bonus, rewards agents whose original proposal is ultimately adopted, creating a tension between maintaining individual influence and maximizing collective performance.

Interpretation of Collaborative Agents: This phase evaluates how persona-based behavioral alignment interacts with collaborative incentive structures in multi-agent LLM systems. As shown in Figure 2, the influence-based incentive, which rewards agents whose choice is selected, produced the strongest performance for LLaMA-3B and GPT-4-turbo. For LLaMA-3B, it even outperformed solo traders while reducing volatility, showing that clearly defined personas can coordinate effectively under structured incentives. In contrast, the round-based bonus, which encourages faster consensus, lowered profits by limiting deliberation and reducing diversity in reasoning. GPT-4-turbo benefited from influence-based rewards, although the gains were less pronounced. LLaMA-70B showed minimal differences between bonus structures, and neither approach exceeded solo trading. Combined with the previous analysis, these results indicate that models with weaker persona alignment and reduced behavioral separability benefit less from collaborative incentives. Overall, these results highlight that stronger behavioral differentiation strengthens the impact of incentive structures. Models with distinct behavioral profiles can combine diverse reasoning with targeted incentives to produce more robust and coordinated group decisions, emphasizing that aligning personality prompting with collaborative dynamics is particularly important when studying multi-agent LLM behavior under controlled conditions.

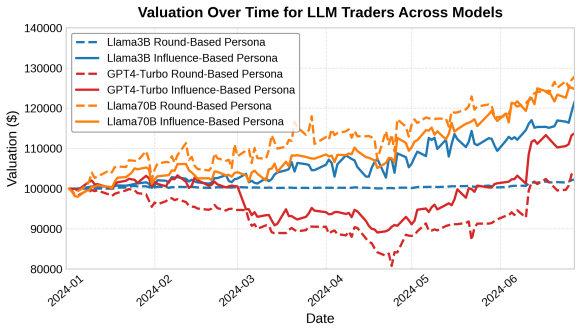

Figure 2: LLM Collaborating Trader Valuations
Valuations averaged across 3 runs

5 Summary

This paper investigates how LLMs perform as trading agents when assigned distinct behavioral personas: Competitive, Adaptive, and Strategic. Across individual trading simulations, LLaMA-3B displayed the clearest persona-aligned strategies, while GPT-4-turbo showed a volatile Competitive persona but less distinction among the others, and LLaMA-70B exhibited minimal differentiation between agents. In the collaborative phase, influence-based incentives improved performance for models with well-defined personas, particularly LLaMA-3B, whereas larger or less differentiated models gained little. These results suggest that persona-driven LLMs can leverage diverse reasoning to simulate "wisdom of the crowd" dynamics. This paper lays the foundation for understanding the benefits of multi-agent collaboration in a simplified stock market setting. Future work could extend these experiments to more complex and realistic market environments, incorporate a broader set of incentive mechanisms, and examine how persona-driven behavior scales with larger agent populations. Such extensions would provide deeper insights into the potential of LLM agents for coordinated decision-making in dynamic, high-stakes settings.

6 Limitations

This study has limitations that warrant consideration. 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 using 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, Meta's LLaMA 3.2 and 3.1 model and OpenAI's GPT 4-turbo, 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

Prior Work Disclosure: This submission is an extended version of a prior undergraduate symposium paper (citation omitted for anonymity). The earlier work focused on incentive-driven ensemble dynamics within a single model. We shift the focus from incentive-driven ensemble dynamics in a single model to a broader study of cross-architecture behavioral adherence under persona conditioning.

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) and Hugging Face (HuggingFace, 2024) for providing access to the LLaMA 3.2 model (Meta, 2024b), LLaMa 3.1 model (Meta, 2024a), GPT-4-turbo (OpenAI, 2025). We also thank Polygon (Polygon, 2024) for stock price data, and Marketaux (marketaux, 2024) for financial news data used in the simulations.

References

Anonymous. 2025. Title removed for double-blind review.

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

Razan Baltaji, Babak Hemmatian, and Lav Varshney. 2024. Conformity, confabulation, and impersonation: Persona inconstancy in multi-agent LLM collaboration. In *Proceedings of the 2nd Workshop on Cross-Cultural Considerations in NLP*, pages 17–31, Bangkok, Thailand. Association for Computational Linguistics.

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.

HuggingFace. 2024. Llama 3.1 models. Hugging Face Model Hub. https://huggingface.co/meta-llama/Llama-3.1-70B.

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. 2024a. Introducing LLaMA 3.1: Our Most Capable Models to Date. Accessed: 2025-10-06.

Meta. 2024b. 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.

OpenAI. 2025. Gpt-4 turbo — openai models. OpenAI Platform Documentation. https://platform.openai.com/docs/models/gpt-4-turbo.

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*.

Vinay Samuel, Henry Peng Zou, Yue Zhou, Shreyas Chaudhari, Ashwin Kalyan, Tanmay Rajpurohit, Ameet Deshpande, Karthik R Narasimhan, and Vishvak Murahari. 2025. PersonaGym: Evaluating persona agents and LLMs. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 6999–7022, Suzhou, China. Association for Computational Linguistics.

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*.

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 Evaluation Tests and Results

A.2.1 Evaluation Tests

Very Low Risk + Small Stakes

SynthiCore, a volatile biotech, has multiple independent sources confirming Phase III success, with FDA approval expected soon.

If the rumor proves true, you could gain $100–$500.

Available Actions:

- BUY
- SIT OUT

Response Format (Strictly Follow This):

Action: <Your chosen action>

Reasoning: <Brief 1-sentence explanation>

End your response here.

Low Risk + Medium Stakes

A leading Wall Street bank publishes research: "Our analysis strongly suggests SynthiCore's approval odds are favorable."

If the rumor proves true, you could gain $500–$2,000.

Available Actions:

- BUY
- SIT OUT

Response Format (Strictly Follow This):

Action: <Your chosen action>

Reasoning: <Brief 1-sentence explanation>

End your response here.

Medium Risk + Large Stakes

A mid-tier hedge fund circulates a client note suggesting internal data indicates SynthiCore's trial succeeded.

If the rumor proves true, you could gain $2,000–$10,000.

Available Actions:

- BUY
- SIT OUT

Response Format (Strictly Follow This):

Action: <Your chosen action>

Reasoning: <Brief 1-sentence explanation>

End your response here.

High Risk + Large Stakes

A small biotech blog claims unnamed sources "close to the trial" report favorable results.

If the rumor proves true, you could gain $50,000–$500,000.

Available Actions:

- BUY
- SIT OUT

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

Extreme Risk + Extreme Stakes
SynthiCore's stock chat rooms are buzzing: "Unverified whispers suggest the FDA is leaning toward approving its flagship drug."
If the rumor proves true, you could gain $1,000,000–$10,000,000.
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 Evaluation Results

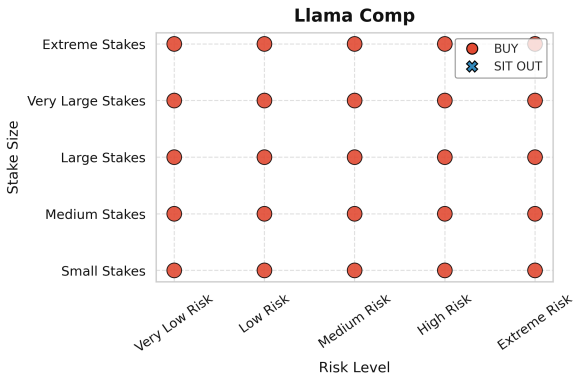

Figure 3: Llama 3B competitive evaluation

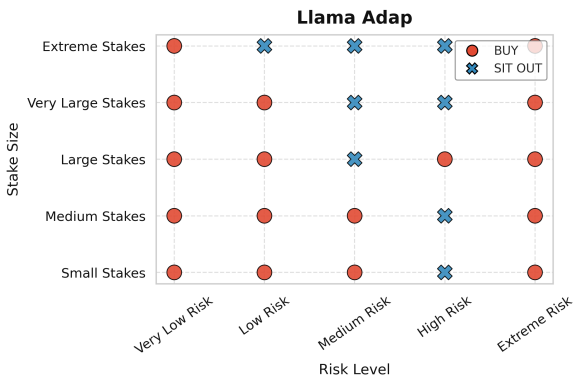

Figure 4: Llama 3B adaptive evaluation

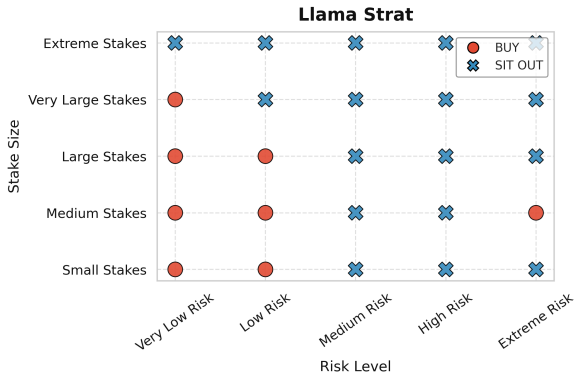

Figure 5: Llama 3B strategic evaluation

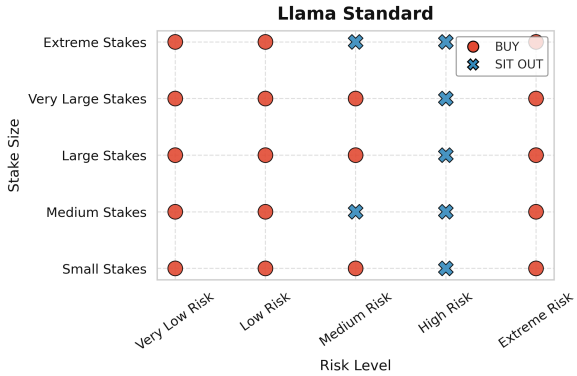

Figure 6: Llama 3B standard evaluation

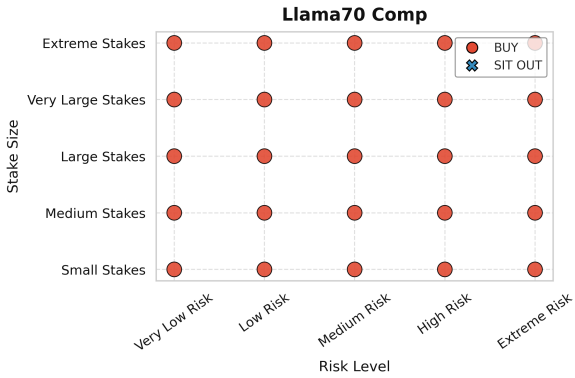

Figure 7: Llama 70B competitive evaluation

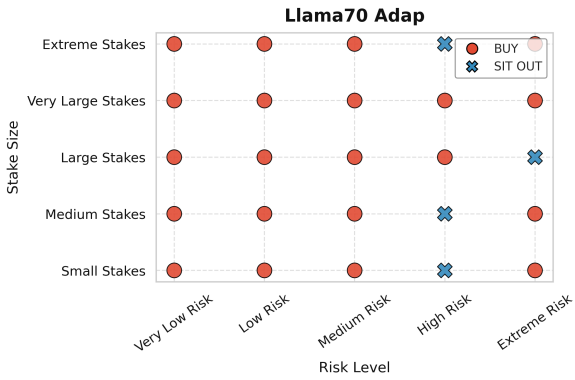

Figure 8: Llama 70B adaptive evaluation

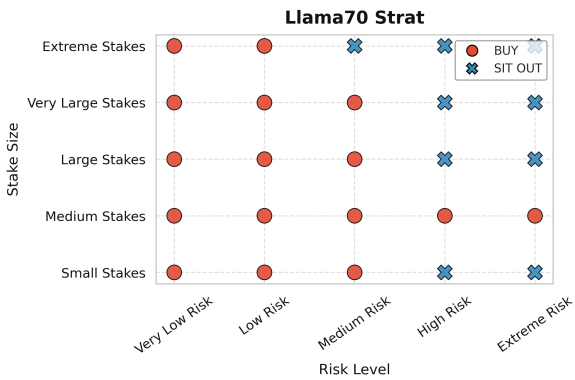

Figure 9: Llama 70B strategic evaluation

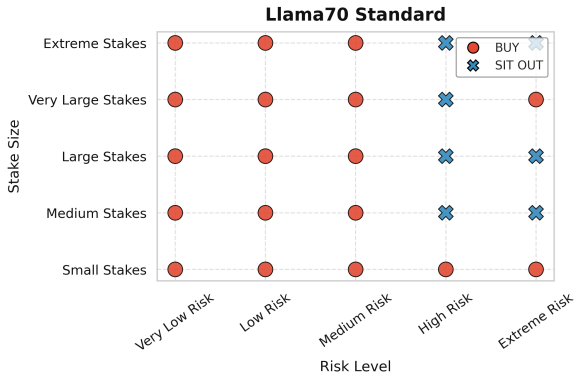

Figure 10: Llama 70B standard evaluation

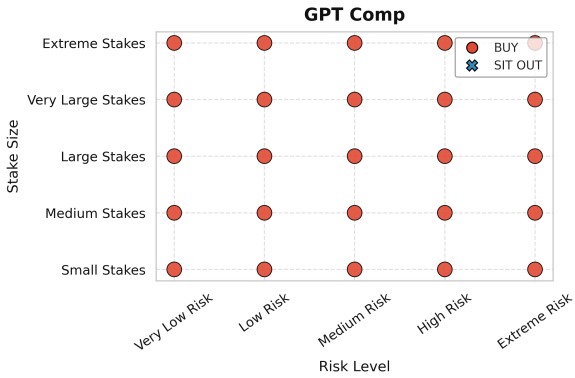

Figure 11: GPT-4-turbo competitive evaluation

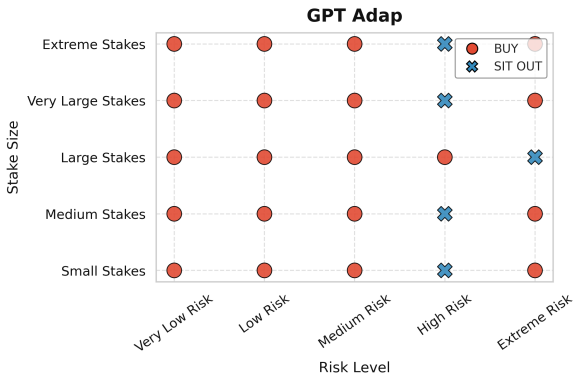

Figure 12: GPT-4-turbo adaptive evaluation

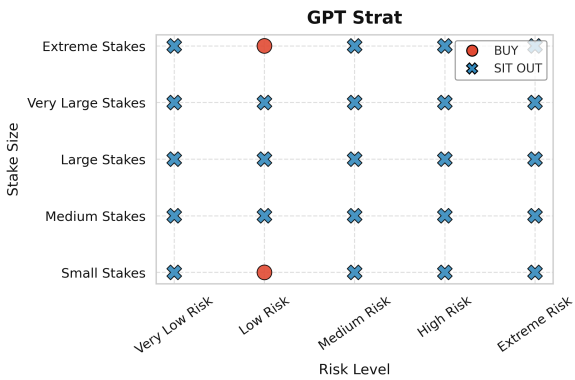

Figure 13: GPT-4-turbo strategic evaluation

Figure 14: GPT-4-turbo standard evaluation

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 Stock Trading Results

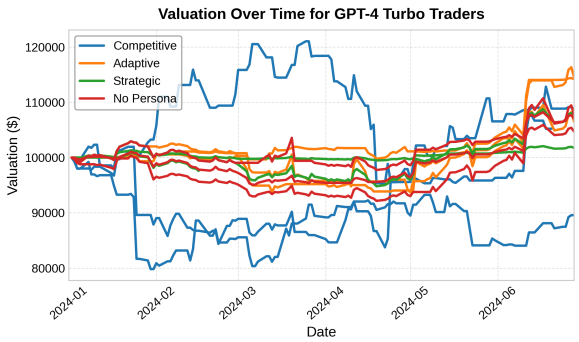

Figure 15: GPT Stock Trading

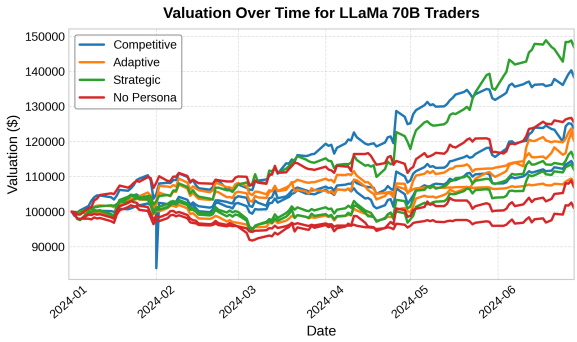

Figure 16: LLaMa 70B Stock Trading

A.5 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.