# LLM-based Agents in Supply Chain Games: The Role of Incomplete Information and Model Heterogeneity

Extended Abstract

Jiuyun Jiang
Harbin Institute of Technology
Harbin, China
jiuyunjiang@stu.hit.edu.cn

Yuecheng Hong
Harbin Institute of Technology
Harbin, China
2023112391@stu.hit.edu.cn

Jiangnan Shi
Harbin Institute of Technology
Harbin, China
jnshi@stu.hit.edu.cn

Song Huang
Harbin Institute of Technology
Harbin, China
23b910018@stu.hit.edu.cn

Jin Yang
Harbin Institute of Technology
Harbin, China
22B910011@stu.hit.edu.cn

Guangxin Jiang
Harbin Institute of Technology
Harbin, China
gxjiang@hit.edu.cn

Xiaomeng Guo
Hong Kong Polytechnic University
Hong Kong, China
xiaomeng.guo@polyu.edu.hk

Guang Xiao
Hong Kong Polytechnic University
Hong Kong, China
guang.xiao@polyu.edu.hk

## ABSTRACT

Effective collaboration is essential for mitigating market volatility, yet complete information sharing among partners is often impractical. By employing diverse Large Language Models as autonomous agents, we design controlled experiments in which information is shared only among subsets of enterprises, approximating realistic business environments. Our results reveal a counterintuitive finding: partial information sharing can generate system level benefits comparable to those achieved under full transparency. We further compare agent behavior and identify differences in decision stability. DeepSeek exhibiting the most consistent performance, followed by Qwen and Llama. Finally, experiments within a Llama based environment show that introducing a higher capability model can improve both stability and aggregate performance. Overall, our study provides a scalable experimental framework for artificial society modeling and demonstrates the potential of LLM-based agent simulations for investigating complex socio economic systems.

## KEYWORDS

Multi-Agent Systems, Large Language Models, Supply Chain Management, Incomplete Information Games, Agent-Based Simulation

**ACM Reference Format:**
Jiuyun Jiang, Yuecheng Hong, Jiangnan Shi, Song Huang, Jin Yang, Guangxin Jiang, Xiaomeng Guo, and Guang Xiao. 2026. LLM-based Agents in Supply Chain Games: The Role of Incomplete Information and Model Heterogeneity: Extended Abstract. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 3 pages. https://doi.org/10.65109/UYIA3362

Guangxin Jiang is the Corresponding Author.



*Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus.* © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). https://doi.org/10.65109/UYIA3362

## 1 INTRODUCTION

Effective collaboration is central to modern socio economic systems, especially in global supply chains where timely information sharing helps mitigate market volatility [5, 11, 21]. In practice, however, inter firm collaboration is constrained by information leakage risks and limited trust, rendering complete information sharing infeasible [2, 18, 26, 27]. As a result, partial information sharing emerges as a pervasive yet insufficiently understood organizational form. Modeling such environments remains challenging, as analytical approaches rely on strong simplifying assumptions and human subject experiments face scalability and cost constraints, limiting their ability to capture dynamic multi stage decision making [12, 16, 28].

Recent advances in Large Language Models offer a new paradigm for agent-based simulation of complex social systems [3, 4, 7, 8, 10, 15, 23]. Prior work demonstrates that LLM-based agents can reproduce human decision patterns and strategic reasoning in classical game theoretic settings [13, 17, 22]. Nevertheless, existing studies largely focus on simple or static games, falling short of the dynamic and repeated nature of complex scenarios [1, 14, 19, 24, 25, 29].

To fill this gap, we develop a multi stage supply chain simulation populated by heterogeneous LLM agents. By deploying diverse agents, we capture heterogeneity in decision capabilities. Extensive multi round experiments allow us to examine how information structures and agent heterogeneity jointly shape collective outcomes in complex socio economic systems.

## 2 PRELIMINARIES

We model the supply chain as a society of firms to study coordination mechanisms, with an emphasis on information sharing (IS). We instantiate this artificial society in the canonical Beer Game framework [6, 9], which provides a standardized and empirically grounded testbed. Each society comprises four agents representing the retailer (S1), wholesaler (S2), distributor (S3), and manufacturer (S4), interacting over $T$ periods. Shipments $S_t^{i,g}$ and inventories $I_t^{i,g}$ evolve according to the Beer Game dynamics, where the primary

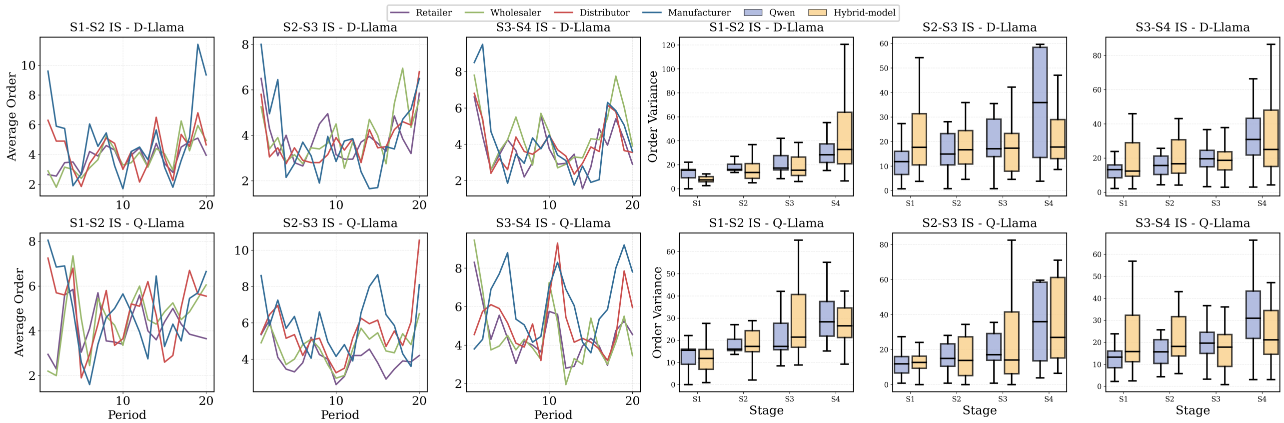


**Figure 1: Performance of hybrid agent teams (D-Llama and Q-Llama) in partial IS scenarios. Left panels show stable order dynamics. Right panels compare order variance of the hybrid teams (blue) against a pure Qwen baseline (orange), revealing highly similar performance profiles.**

decision is the order quantity $O_t^{i,g}$:

$$S_t^{i,g} = \begin{cases} \min\left\{D_t, \max\left[I_{t-1}^{i,g} + S_{t-2}^{i+1,g}, 0\right]\right\} & \text{for } i = 1,\\ \min\left\{O_{t-2}^{i-1,g}, \max\left[I_{t-1}^{i,g} + S_{t-2}^{i+1,g}, 0\right]\right\} & \text{for } i = 2,3,\\ \min\left\{O_{t-2}^{i-1,g}, \max\left[I_{t-1}^{i,g} + O_{t-3}^{i,g}, 0\right]\right\} & \text{for } i = 4,\end{cases}$$

with inventory updates, where negative inventory denotes backlog:

$$I_t^{i,g} = \begin{cases} I_{t-1}^{i,g} + S_{t-2}^{i+1,g} - D_t & \text{for } i = 1,\\ I_{t-1}^{i,g} + S_{t-2}^{i+1,g} - O_{t-2}^{i-1,g} & \text{for } i = 2,3,\\ I_{t-1}^{i,g} + O_{t-3}^{i,g} - O_{t-2}^{i-1,g} & \text{for } i = 4.\end{cases}$$

Each agent minimizes its individual total cost, defined as inventory holding cost plus backlog cost.

We populate societies with open source LLM agents from distinct model families and scales, namely **DeepSeek-V3.1**, **Qwen 2.5-32B**, and **Llama 3.1-8B**. All agents operate with temperature 1 for consistency, following prior large scale scenario experiments [7]. Our design varies both agent composition and information flow through two experiment classes. First, in **homogeneous societies**, all four agents use the same model and we evaluate five information regimes: **No IS**, **Full IS**, and three localized **Partial IS** alliances (**S1-S2 IS**, **S2-S3 IS**, **S3-S4 IS**). Second, in **heterogeneous societies**, we conduct ensemble experiments by starting from an all Llama 3.1 baseline and replacing agents involved in a partial IS pact with higher capability models to assess system level effects.

Each of the six configurations is **independently replicated 32 times** over 20 periods, consistent with previous study [9, 20]. All agents use Chain-of-Thought (CoT) prompting to support structured and transparent decision-making.

## 3 EXPERIMENT

First, we validate our LLM-based simulation paradigm by demonstrating its ability to replicate foundational behavioral patterns observed in human experiments. Statistical analyses based on the Mann–Whitney U test confirm that information sharing significantly reduces order variance under all configurations, consistent with prior experimental evidence [6]. The results but only validate the effectiveness of full information sharing but prove that near optimal coordination does not require perfect information. Then we investigate if the choice of LLM is merely an implementation detail or a critical factor that dictates the behavior of the entire multi agent system. We performed one tailed Mann Whitney U tests on the total order variance across all five information sharing conditions between diverse models. As demonstrated in Table 1, DeepSeek demonstrates significantly lower variance than both Qwen and Llama across every scenario. Similarly, Qwen exhibits significantly lower variance than Llama in nearly all conditions.

**Table 1: Pairwise Model Stability Comparison Using One-Tailed Mann-Whitney U Test.**

| **Models** | **Experimental Conditions** | | | | |
|---|---|---|---|---|---|
| | **No IS** | **Full IS** | **S1-S2 IS** | **S2-S3 IS** | **S3-S4 IS** |
| DeepSeek vs Llama | <0.001 | <0.001 | <0.001 | <0.001 | <0.001 |
| DeepSeek vs Qwen | 0.002 | <0.001 | <0.001 | 0.002 | 0.002 |
| Qwen vs Llama | 0.013 | 0.007 | 0.026 | 0.045 | 0.004 |

Furthermore, real world organizations are typically composed of individuals with heterogeneous capabilities. To examine whether the performance of a less stable system be substantially improved by introducing a small number of more capable agents, we construct two hybrid societies, **D-Llama** and **Q-Llama**, by replacing two key agents in a Llama based society with higher capability agents respectively. As illustrated in Figure 1, both configurations exhibit improved operational stability. The time series plots indicate reduced fluctuations, while the box plot comparisons with the homogeneous Qwen baseline reveal a counterintuitive result. The variance distributions of the hybrid societies closely match those of the substantially more capable, fully Qwen based society.

## ACKNOWLEDGMENTS

This work was partially supported by the National Natural Science Foundation of China (No. 72293562 and 724B2012).


## REFERENCES


[1] Elif Akata, Lion Schulz, Julian Coda-Forno, Seong Joon Oh, Matthias Bethge, and Eric Schulz. 2025. Playing repeated games with large language models. *Nature Human Behaviour* (2025), 1–11.

[2] Mohammad M Ali, Mohamed Zied Babai, John E Boylan, and Aris A Syntetos. 2017. Supply chain forecasting when information is not shared. *European Journal of Operational Research* 260, 3 (2017), 984–994.

[3] Lisa P Argyle, Ethan C Busby, Nancy Fulda, Joshua R Gubler, Christopher Rytting, and David Wingate. 2023. Out of one, many: Using language models to simulate human samples. *Political Analysis* 31, 3 (2023), 337–351.

[4] Marcel Binz and Eric Schulz. 2023. Using cognitive psychology to understand GPT-3. *Proceedings of the National Academy of Sciences* 120, 6 (2023), e2218523120.

[5] Dean C Chatfield, Jeon G Kim, Terry P Harrison, and Jack C Hayya. 2004. The bullwhip effect—impact of stochastic lead time, information quality, and information sharing: a simulation study. *Production and operations management* 13, 4 (2004), 340–353.

[6] Rachel Croson and Karen Donohue. 2006. Behavioral causes of the bullwhip effect and the observed value of inventory information. *Management Science* 52, 3 (2006), 323–336.

[7] Ziyan Cui, Ning Li, and Huaikang Zhou. 2025. A large-scale replication of scenario-based experiments in psychology and management using large language models. *Nature Computational Science* (2025), 1–8.

[8] Ishita Dasgupta, Andrew K Lampinen, Stephanie CY Chan, Hannah R Sheahan, Antonia Creswell, Dharshan Kumaran, James L McClelland, and Felix Hill. 2022. Language models show human-like content effects on reasoning tasks. *arXiv preprint arXiv:2207.07051* (2022).

[9] Andrew M Davis, Blair Flicker, Kyle Hyndman, Elena Katok, Samantha Keppler, Stephen Leider, Xiaoyang Long, and Jordan D Tong. 2023. A replication study of operations management experiments in Management Science. *Management Science* 69, 9 (2023), 4977–4991.

[10] Andrew M Davis, Shawn Mankad, Charles J Corbett, and Elena Katok. 2024. OM Forum—The best of both worlds: Machine learning and behavioral science in operations management. *Manufacturing & Service Operations Management* 26, 5 (2024), 1605–1621.

[11] Jeroen Dejonckheere, Stephen M Disney, Marc R Lambrecht, and Denis R Towill. 2004. The impact of information enrichment on the bullwhip effect in supply chains: A control engineering perspective. *European journal of operational research* 153, 3 (2004), 727–750.

[12] Roberto Dominguez, Salvatore Cannella, Ana P Barbosa-Póvoa, and Jose M Framinan. 2018. OVAP: A strategy to implement partial information sharing among supply chain retailers. *Transportation Research Part E: Logistics and Transportation Review* 110 (2018), 122–136.

[13] Jinhao Duan, Renming Zhang, James Diffenderfer, Bhavya Kailkhura, Lichao Sun, Elias Stengel-Eskin, Mohit Bansal, Tianlong Chen, and Kaidi Xu. 2024. Gtbench: Uncovering the strategic reasoning limitations of LLMs via game-theoretic evaluations. *arXiv preprint arXiv:2402.12348* (2024).

[14] Fulin Guo. 2023. GPT in game theory experiments. *arXiv preprint arXiv:2305.05516* (2023).

[15] Thilo Hagendorff, Sarah Fabi, and Michal Kosinski. 2023. Human-like intuitive behavior and reasoning biases emerged in large language models but disappeared in ChatGPT. *Nature Computational Science* 3, 10 (2023), 833–838.

[16] Jan Holmström, Johanna Småros, Stephen M Disney, and Denis R Towill. 2016. Collaborative supply chain configurations: The implications for supplier performance in production and inventory control. In *Developments in Logistics and Supply Chain Management: Past, Present and Future*. Springer, 27–37.

[17] John J Horton. 2023. *Large language models as simulated economic agents: What can we learn from homo silicus?* Technical Report. National Bureau of Economic Research.

[18] Yeu-Shiang Huang, Ming-Chi Li, and Jyh-Wen Ho. 2016. Determination of the optimal degree of information sharing in a two-echelon supply chain. *International Journal of Production Research* 54, 5 (2016), 1518–1534.

[19] Ilya Jackson, Dmitry Ivanov, Alexandre Dolgui, and Jafar Namdar. 2024. Generative artificial intelligence in supply chain and operations management: a capability-based framework for analysis and implementation. *International Journal of Production Research* 62, 17 (2024), 6120–6145.

[20] Samuel Kirshner. 2024. Artificial agents and operations management decision-making. *UNSW Business School Research Paper Forthcoming* (2024).

[21] Hau L Lee, Venkata Padmanabhan, and Seungjin Whang. 1997. Information distortion in a supply chain: The bullwhip effect. *Management science* 43, 4 (1997), 546–558.

[22] Juanjuan Meng. 2024. AI emerges as the frontier in behavioral science. *Proceedings of the National Academy of Sciences* 121, 10 (2024), e2401336121.

[23] 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*. 1–22.

[24] Steve Phelps and Yvan I Russell. 2023. Investigating emergent goal-like behaviour in large language models using experimental economics. *arXiv preprint arXiv:2305.07970* (2023).

[25] Yinzhu Quan and Zefang Liu. 2024. Invagent: A large language model based multi-agent system for inventory management in supply chains. *arXiv preprint arXiv:2407.11384* (2024).

[26] Matan Shnaiderman and Fouad El Ouardighi. 2014. The impact of partial information sharing in a two-echelon supply chain. *Operations Research Letters* 42, 3 (2014), 234–237.

[27] Robert Spekman and Edward W Davis. 2016. The extended enterprise: a decade later. *International Journal of Physical Distribution & Logistics Management* 46, 1 (2016), 43–61.

[28] John D Sterman. 1989. Modeling managerial behavior: Misperceptions of feedback in a dynamic decision making experiment. *Management Science* 35, 3 (1989), 321–339.

[29] Haojie Wang, Jiuyun Jiang, L Jeff Hong, and Guangxin Jiang. 2025. LLMs for Supply Chain Management. *arXiv preprint arXiv:2505.18597* (2025).