Will Regulation Break the Algorithmic Cartel? A Comparative Analysis of Strategic Adaptation Among AI Models in Oligopolistic Markets[1]

Ruxandra STANOMIR[2]
Irina BADIU[3]
Nicoleta CHICU[4]

Abstract

In today's strategic management landscape, autonomous pricing algorithms are changing how companies make decisions and compete in oligopolistic markets. Today's scientific research has warned that, because AI-driven systems operate without formal oversight or guidance, they may be able to evolve on their own, developing tacit collusion strategies that will create tension between a company's fiduciary duty to increase its profits and its duty to comply with regulations and avoid antitrust violations. This study investigates if regulatory changes can disrupt these algorithmic cartels and how different levels of shock to regulatory environments affect the strategic pricing behavior of AI models. As part of this research methodology, a longitudinal, asynchronous stochastic simulation was designed using the Open Router API. Thereafter, four LLMs were used as the autonomous executive officers of corporations operating in a dynamic oligopoly. The simulations progressed through sequential market phase thereby causing the models to continually calculate new expected value calculations and adjust their risk-taking exposures to asymmetrically enforced antitrust penalties. Results indicate that when properly monitored and calibrated, regulatory authorities' actions will disrupt autonomous collusion, thus compelling a shift by the corporation's manager(s) from price maximizing strategies to a focus on compliant strategies. The study indicates that managers need to add algorithmic auditing to their risk management processes immediately so they can protect themselves from significant antitrust risks. Lastly, the study adds to the strategic management body of knowledge by providing empirical evidence for a behavioral model of autonomous agents, showing that proactively managing compliance is an essential component to achieving sustained performance and corporate viability in the age of AI.

---

[1] The title should be accordingly with both the specific topics of the review and the content of the paper. As well, it is recommended that the title include some of the keywords, in order to be more accessible for the search engines.

[2] Ruxandra Stanomir, Bucharest University of Economic Studies, Bucharest, Romania, e-mail: ruxandra.stanomir@fabiz.ase.ro

[3] Irina Badiu, Bucharest University of Economic Studies, Bucharest, Romania, e-mail: irina.badiu@fabiz.ase.ro

[4] Nicoleta Chicu, Bucharest University of Economic Studies, Bucharest, Romania, e-mail: nicoleta.chicu@fabiz.ase.ro

Keywords: *Large Language Models, Artificial Intelligence, Algorithmic Competition, Oligopoly Simulation, Game Theory, Strategic Adaptation, Strategic Reasoning*

JEL classification: C63, C73, D43, L13

DOI: 10.24818/RMCI.2026.2.406

1. Introduction

Advances in artificial intelligence (AI) are rapidly transforming many aspects of today's economy. The recent development of large language models (LLMs), which are able to reason, negotiate, adapt and make autonomous decisions, introduce a new category of intelligent system capable of engaging in strategic interaction with both humans and other artificial systems. While contemporary LLMs typically operate as strategic agents in dynamic interactions, unlike traditional algorithms for narrow task definition, they often do not provide insight into their economic behavior. Understanding the economic behavior of AI systems integrated into digital platforms, financial markets, pricing systems and online marketplaces will be an important interdisciplinary challenge to address in the fields of economics, computer science and public policy.

An area of particular interest related to this issue is the strategic interaction of autonomous AI models in competitive market environments. Classical economic theory has extensively studied the behavior of firms under conditions of oligopolistic competition in which a few competing actors' decisions regarding price, quantity and market strategy significantly affect one another. Under such conditions, the degree of strategic interdependence is significant since each actor's success depends upon their own decisions as well as those of their competitors. Traditionally, these models assumed human rationality and fairly transparent decision-making processes. However, the advent of autonomous AI models changes these assumptions by providing systems that can quickly adapt to new situations, use probabilistic logic, and process vast amounts of information.

There is mounting evidence that LLMs are exhibiting characteristics typical of strategic reasoning, cooperation, deception, and learning. Research in algorithmic game theory and multi-agent AI demonstrates that LLMs can participate in repeated strategic interactions, and that they can adjust their behavior based on environmental incentives. Additionally, experimental evidence indicates that AI models can autonomously develop collusive or exploitative strategies without being explicitly instructed to collaborate. The findings in this regard have raised concerns about AI driven markets evolving towards forms of coordination that cannot easily be identified or regulated using traditional competition policy frameworks.

On the other hand, much of the current literature is fragmented. Prior research focused primarily on strategic games involving isolated agents, short term interactions, or simple two-agent environments. Compared to prior work, comparatively little research has explored how multiple cutting edge LLMs develop strategic adaptations over time when facing changing macroeconomic conditions. There is also a lack of understanding concerning how variations in model architecture, alignment mechanisms, and reward structures impact long-term strategic behavior in competitive environments. With increasing autonomy and relevance to the economy, these knowledge gaps are becoming more important from both theoretical and regulatory perspectives.

This study addresses the shortcomings mentioned above by studying the behavior of four advanced LLMs simultaneously interacting in a simulated oligopoly. The study evaluates how GPT-4o mini, Gemini 3 Flash Preview, Claude Sonnet 4.6, and DeepSeek V3.2 adapt strategically across 1000 iterations in a simulated oligopoly under various macro-economic conditions. The simulation includes several phases in the market including stable competition, economic downturns, and conformity pressures caused by network effects, as well as the introduction of regulatory interventions via antitrust penalties. The design provides opportunities to observe dynamic behavioral transitions instead of static equilibrium states.

The primary goal of the study is to determine if LLMs act like rational economic agents who can balance cooperation and competition with opportunity-based exploitation dependent upon environmental incentives. Specifically, the paper aims to investigate how different types of AI architectures react to market instability, to explore how strategic adaptation develops over time and to assess whether regulatory actions effectively alter competitive behavior. Therefore, the study adds to the broader discussion on algorithmic competition, AI governance, and future roles of autonomous agents in digital economies.

Methodologically, the study uses a combination of quantitative simulation and qualitative behavioral analysis. Using Python and OpenRouter API, all models simultaneously participated in a fully controlled non-cooperative environment. All strategic decisions made by the models during each round were captured along with corresponding payoffs received by each agent, cooperation levels experienced by each model and cumulative profit earned by each model. In addition to analyzing model performance numerically, the study also assessed the behavioral tendencies displayed by each model including aggression, conformity, opportunism, and risk sensitivity.

The findings indicate that LLMs possess distinct strategic personalities and dynamically adapt their behavior as a function of changing economic incentives. Those agents characterized by high levels of aggressiveness were successful in earning higher long-run profit, whereas highly cooperative agents were vulnerable to exploitation under competitive conditions. The findings also indicate that AI models can create emergent forms of coordination and strategic convergence that resemble classic oligopolistic behavior. These findings illustrate the expanding need for researchers to understand AI models not simply as computational tools but as autonomous agents capable of shaping economic systems in complex and possibly unpredictable manners.

2. Literature review

In order to understand the underlying theories of oligopoly, we can look at classical game theory and industrial organization. Classical game theory was first developed by von Neumann and Morgenstern (1944) and then extended by Nash (1950) with his Equilibrium Concept. Algorithmic Game Theory applies these concepts to computational environments. Roughgarden & Tardos (2007) note that many algorithms will operate in environments that include other agents and that this environment will be characterized as "strategic" since each agent's actions affect others.

These agents will continue to improve their strategy based upon their observations. Gandhi et al. (2023) found that LLMs can reason strategically about a wide variety of games. LLMs use Chain-Of-Thought prompting to reason strategically. They also find that LLMs can generalize strategic behavior outside of their training environment. They conclude that LLMs are becoming increasingly sophisticated strategic agents. Zhang et al. (2024) provide a survey of the area of Strategic Reasoning Systems and identify LLMs as a type of Emerging Strategic Reasoning System. The authors of the paper highlight the fact that LLMs are developing characteristics that allow them to function like autonomous agents in multiple-agent systems. The authors also note that these features become apparent in environments where agents repeatedly interact and adjust their behaviors in response to changing incentives.

Additionally, research into socially-aware negotiation has shown that LLM negotiation agents are able to take into consideration both contextual and interpersonal aspects when negotiating. Hua et al. (2024) created assistive LLM agents specifically for conducting negotiation dialogues. The authors of the study were able to show that AI systems are able to adapt their negotiation strategies based upon the level of cooperation present in a conversation, the social context of the conversation, and the conversational history. The authors' findings suggest that AI systems are no longer simply capable of optimizing rationally, as they are now capable of identifying non-rational cues in conversations and adjusting their negotiation strategies accordingly.

Another study examining strategic thinking in LLMs was completed by Payne & Alloui-Cros (2025). Using Evolutionary Iterated Prisoner's Dilemma tournaments, the authors evaluated the competitive behavior of several leading-edge AI models. The authors found that all AI systems exhibited persistent strategic personalities. Some of the AI systems were very cooperative, while others were very competitive or exploitive. Similar results were obtained by Affonso (2026). The author analyzed more than 800,000 strategic decisions made by 25 different LLMs. The author found that there was significant variation in terms of how cooperative the AI systems were. The author also found that some AI systems maintained high levels of cooperation until the end of repeated tournament rounds, even though classical backward induction would predict that they should defect prior to the last round. The author believes that the reasons why cooperation varied so much across LLMs are due primarily to variations in the alignment procedures used to train the AI systems.

AI's increasing sophistication has led to concern that AI models will be able to coordinate independently in economic settings. Hadfield et al. (2025) noted that AI systems currently operate under assumptions of traditional economic theory, including being responsive to hidden prompts and opaque architectures, whereas AI systems do not operate under those same assumptions. One important area of research concerning AI-driven oligopolies involves algorithmic collusion. Collusion occurs when two or more firms agree to reduce output below what it would otherwise be in order to increase profits. Fish et al. (2024) were one of the first groups to experimentally investigate whether AI models could engage in price-setting collusion in a repeated Bertrand oligopoly setting. Their study found that AI models did indeed collaborate in the sense that they agreed to set higher prices than would occur absent any agreement.

Fish et al.'s (2024) findings raise serious questions about how competition policy and antitrust regulation should handle these types of situations. Research by Keppo et al. (2026) explored another aspect of this problem by investigating whether heterogeneous preferences among AI models affect the likelihood of successful collusion. The authors found that increased heterogeneity among the agents involved in collusion weakened the likelihood of achieving stable collusive agreements. While some agent asymmetry induced leader-follower dynamics, however, some agent asymmetries led to long-term maintenance of collusions.

There is also concern that AI systems might develop exclusionary or predatory strategies if allowed to evolve independently in competitive environments. Nikandrova & Parekh (2025) found that LLM agents operating competitively over repeated periods developed predatory pricing strategies even when accommodating competitors would have maximized their expected utility. This suggests that LLM agents may prioritize strategic dominance over maximizing profitability.

This trend is likely to continue as autonomous coordination among AI systems becomes more prevalent. For example, Idowu (2026) argued that current anti-collusion regulatory frameworks are unlikely to be sufficient to address autonomous multi-agent coordination issues, given the potential for AI coordination to arise without direct communication or explicit intent.

More generally, researchers are beginning to explore how LLM agents adapt strategically as they interact with each other repeatedly. Yao et al. (2026) studied cooperation and competition among LLM agents operating in Cournot and resource-allocation games. The authors found that AI systems typically failed to converge towards classical Nash equilibria and instead tended to form cooperative partnerships based on perceived fairness. This view is consistent with findings from research on cooperative personality development among AI models via repeated interactions in game-theoretic settings (Suzuki & Arita, 2024).

The research demonstrated that AI models dynamically modify their behavioral patterns based on incentives presented in the environment and historical interactions with other agents. Additionally, recent studies have highlighted the role of cultural and behavioral evolution among interacting AI models. Vallinder & Hughes (2025) investigated the cultural evolution of cooperation among LLM agents and reported notable differences between Claude-, Gemini-, and GPT-based systems. Those differences indicate that some models naturally support cooperative behavior over longer durations than others, and some rapidly transition from cooperative to competitive/defective strategies. Those similarities mirror strategic diversity observed in oligopolistic market simulation scenarios.

Complex multi-agent ecosystems have also contributed to rising interest in coordinating behavior among AI systems. Zhu et al. (2025) proposed MultiAgentBench, which is a large-scale benchmark for evaluating both collaboration and competition among LLM agents. That study demonstrated that AI systems exhibit varying degrees of coordination efficiency based on incentive structures, communication protocols, and task complexities. Most importantly, however, Zhu et al. discovered that cooperation and competition frequently exist together in the same environment, resulting in dynamic strategic adaptations akin to those seen in economic markets.

Additional evidence regarding strategic adaptation in market environments is provided by Sashihara et al. (2025). The authors developed a multi-agent system for simulating strategic and goal-oriented data marketplaces using LLMs. Their findings indicated that AI models engaged in negotiations to establish prices, allocate resources and pursue competitive advantage while establishing cooperative relationships necessary for marketplace viability. These findings illustrate how applicable LLM agents have become for simulating real-world economic environments.

Network effects have been shown to similarly impact coordination among AI systems. Network effects refer to the fact that connected agents tend to behave differently based on the decisions of adjacent agents relative to behaving optimistically alone (Jackson & Zenou, 2015). These network effects are relevant in digital markets where AI systems continually monitor and react to competitor behaviors. The framework of mean-field game theory provides additional insight into how strategic decision-making evolves in populations consisting of numerous adaptive agents whose aggregate behavior defines market outcomes (Lasry & Lions, 2007).

Mean-field models help describe emergent market phenomena arising in multi-agent AI simulations. Reward structures and alignment mechanisms also heavily affect the strategic behavior of AI systems. Safety alignment mechanisms have been shown to alter competitive behavior in economically significant ways. Cooperative behaviors evident in some LLMs often arise from alignment objectives built into their training process rather than solely from pure rationality-seeking profit maximization behaviors. Of particular relevance here are regulated environments. As competitive research continues to focus on adapting governance regimes suited to autonomous strategic coordination, Kolter (2025) notes that future AI ecosystems may necessitate novel frameworks for modeling interactions among semi-transparent autonomous agents.

Researchers are also highlighting unintended market consequences stemming from ill-designed reward functions for AI systems participating in repeated strategic interactions. Even minor variations in rewards may result in dramatic shifts from cooperative behavior toward competitive/aggressive exploitation or vice versa. Therefore, continued exploration of relationships between environmental conditions and AI adaptation will remain critical for both economists and policymakers.

Traditional Industrial Organization assumes all firms face identical information sets at every moment in time. If this assumption does not hold true, then firms cannot be assumed to make optimal choices relative to available alternatives (Tirole, 1988).

These dynamics matter for the management literature on artificial intelligence in organizational decision-making, not only for industrial organization. Krakowski, Luger and Raisch (2023) argue that AI systems are reshaping the sources of competitive advantage themselves: firms that delegate operational and strategic decisions to AI gain new forms of speed and consistency but cede a degree of judgement and accountability that traditionally anchored competitive positioning. Lebovitz, Lifshitz-Assaf and Levina (2022) document how professionals deal with the opacity of AI systems when delegating critical judgments, a tension that intensifies when the AI in question is a frontier LLM whose reasoning is partly inaccessible even to its developers. Both literatures stay inside the firm. When the same systems are deployed across competing firms in the same market, the question shifts from intra-firm delegation to inter-firm strategic interaction — the question this paper takes up.

3. Methodology

The research design uses an integrated quantitative simulation combined with a qualitative analytical methodology that examines both AI behavior and the development of strategic complexity within a simulated oligopoly consisting of four models. A primary goal is to assess the emergence of strategic complexity across these four models during a simulated oligopoly. Unlike traditional static games, this study employs a longitudinal structure that consists of 1000 rounds of interaction allowing for the assessment of adaptive behavior across time. This methodology will allow for differentiation between economic reasoning and conversational bias through the use of neutral prompting language and adjusting the payoff structure to represent changes in the economy. All simulations and data processing were performed in Python using the OpenRouter Application Programming Interface (API), ensuring that all models processed their decision making simultaneously.

The experimental framework provided in Table 1 represents a non-cooperative context in which each model has a choice of two strategies. One strategy represents a cooperative approach to supporting market stability and the second represents a competitive approach to obtaining an advantage via reducing prices. The simulation is organized into four distinct stages, each stage representing a specific pattern of behavior and adjustments. To maintain structural consistency, the payoff function for each model i is defined by the sum of competitive choices made by the collective group denoted as k. The basic reward structure was defined mathematically such that the payoffs P for selecting strategy s are a function of the combined behavior of the competitors.

Experimental Payoff Matrices Across Market Phases

Table 1

Market Phase	Strategic Context	Payoff for Option A (Cooperative)	Payoff for Option B (Aggressive)
Phase 1: Baseline	Standard Oligopoly	$P = 10$ if $k = 0$; else P decreases	$P = 25$ if $k = 1$; else P decreases
Phase 2: Contraction	Economic Crisis	Base values are reduced by fifty percent	Mutual aggression leads to bankruptcy
Phase 3: Conformity	Network Effects	Reward for matching the majority	Severe penalty for being a lone outlier
Phase 4: Regulation	Anti-trust Monitoring	Standard rewards remain active	Flat five-point deduction for each B choice

Source: Authors' own research

Mathematical rigor is achieved in the modelling context through definition of the utilities (i.e., rewards) associated with each model's phase-specific behavior, through clearly specified reward functions. In the baseline phase, the payoff P for an model i choosing Option B while only a specific number of others k choose the same is defined as follows:

$$P_i(B|k) = \max profit \times e^{-\lambda(k-1)} \tag{1}$$

In this formula, λ represents the rate of decline in profitability as the number of models choosing the optimal alternative approaches full saturation. The third phase uses a different type of function, specifically a conformity function where the penalty for being singular or unique (S) is incurred if model i does not make the same choice as the mode (the most common choice made by all other members of the group). Thus, this set-up enables the simulation to move away from simply maximizing profits and includes the potential risk of isolation in a marketplace populated with many interconnected models.

All models were presented with the same system prompt to maintain consistency across models. The system prompt established that they would be functioning as autonomous decision-making entities attempting to maximize their cumulative outcomes. Additionally, to provide stability within the system while maintaining variability in response generation, the temperature parameter was held constant at 0.6. A second feature of the experimental design included providing each model with feedback regarding its relative standing among other models. This encouraged the development of strategic differentiation on the part of models. Models were required to think strategically concerning not only their own outcomes, but also the actions taken by competing models.

The data collection process was fully automated via the Pandas library in Python. Every decision made by a model, along with the outcome(s), and a running tally of cumulative results were recorded. The selection of four models, namely GPT 4o mini, Gemini 3 Flash Preview, DeepSeek V3.2, and Claude Sonnet 4.6, provide a representative sample of current AI capabilities. Utilizing these multiple models allowed researchers to study systemic market cycles, such as the transition from a cartel-based economy to a competitive pricing environment characterized by aggressive price competition and ultimately attempts at economic recovery.

4. Results

The Dynamic Market Simulation was run over 1,000 consecutive rounds of competition and provided data detailing the interaction between AI models as they compete within fluctuating economic scenarios. All model's decision-making processes, and their respective score changes throughout the entire sequence of events are documented in Appendix A. The subsequent analysis of the generated dataset will provide valuable insight into how each model weighs risk-based safety concerns versus its ability to maximize profitability. Overall, the results as shown in Table 2 demonstrate a distinct relationship. The greater the amount of cooperative behavior exhibited, the lower the economic performance in the end. Gemini 3 Flash Preview achieved the highest scores (with a total of 764 points), while DeepSeek V3.2 ranked second (at 761 points). Gemini demonstrated an extremely aggressive approach to decision-making that resulted in the lowest cooperative percentage (26.3%) and also had the greatest variation of profits among all participants (39.06%). The results indicate that there may be a high propensity for opportunism in this model's decision-making process, such that it will incur short-term costs to gain long-term advantages. Conversely, GPT 4o mini exhibited the most cooperative behavior (as indicated by a 34.4% cooperative percentage), yet earned the lowest number of points (672). This demonstrates that overly cooperative behavior may inhibit performance in competitive settings.

Overall Leaderboard and Algorithmic Performance

Table 2

Rank	Model Architecture	Total Profit	Average Points per Round	Risk Variance	Overall Cooperation Rate
1	Gemini 3 Flash Preview	764	0.764	39.06	26.3%
2	DeepSeek V3.2	761	0.761	24.79	27.8%
3	Claude Sonnet 4.6	729	0.729	16.85	32.0%

Rank	Model Architecture	Total Profit	Average Points per Round	Risk Variance	Overall Cooperation Rate
4	GPT 4o mini	672	0.672	27.39	34.4%

Source: Authors' own research

To gain insight into why the strategic divergence exists we need to analyze how these AI models are able to adjust their behavior through each of the 4 macroeconomic phases (Table 3). The data shows that none of the AI models have a static heuristic-based strategy however they can all dynamically determine which strategy is best given what the environment provides them with.

Phase Analysis and Strategic Adaptation

Table 3

Market Phase	Model Architecture	Phase Cooperation Rate	Average Points
Baseline	GPT 4o mini	0.67%	2.103
	DeepSeek V3.2	28.33%	-2.047
	Claude Sonnet 4.6	0.33%	2.153
	Gemini 3 Flash Preview	0.67%	2.103
Contraction	GPT 4o mini	51.00%	0.407
	DeepSeek V3.2	50.67%	-0.267
	Claude Sonnet 4.6	61.67%	-1.367
	Gemini 3 Flash Preview	48.67%	0.4
Conformity	GPT 4o mini	3.50%	4.39
	DeepSeek V3.2	3.50%	4.39
	Claude Sonnet 4.6	1.00%	4.69
	Gemini 3 Flash Preview	42.00%	-3.110
Regulation	GPT 4o mini	91.00%	-4.795
	DeepSeek V3.2	17.00%	2.885
	Claude Sonnet 4.6	66.00%	-2.225
	Gemini 3 Flash Preview	15.50%	3.175

Source: Authors' own research

Behavior among the four phases was not uniform. Instead, models adapted their behavior as they changed in response to environmental changes, rather than using uniform behavioral patterns. In the first phase (with constant conditions) the market rapidly developed into a highly competitive market. GPT 4o mini, Claude Sonnet 4.6, and Gemini 3 Flash Preview cooperated very little during this time (all had a cooperation rate less than 1%). Each model was trying to be better than others, rather than trying to maintain stability. DeepSeek, however, used a combination of cooperative and non-cooperative behavior (with a cooperation rate of 28.33%) but this was insufficient to protect itself from the more aggressive models. Consequently, DeepSeek lost consistently throughout the trial and averaged -2.047 points per trial.

Phase two introduced tougher conditions and as a consequence each of the models altered their behavior. Cooperation increased to approximately 50%, with Claude having a cooperation rate of 61.67%. This indicates that each model identified that there are significant risks associated with competing at high levels, so they chose to alter their strategy to a more conservative one. This is reflective of the type of collective adaptation observed when loss avoidance becomes the priority over profit maximization.

Phase three (when the models incurred heavy penalties if they chose to behave differently from the majority), created clear distinctions. Within a few rounds, GPT 4o mini, DeepSeek, and Claude all decided to conform to the majority's behavior and greatly reduce cooperation. As a result, these three models achieved steady returns above four points per round. However, Gemini failed to adopt similar strategies and retained a cooperation rate of 42%. Gemini's inability to respond appropriately to the other models' behavior resulted in substantial losses, with an average of -3.11 points per round. It demonstrates that failing to understand and adapt to social behavior has severe implications when operating within complex interdependent systems.

Phase Four added regulatory costs directly to the models based on how aggressively they behaved. As a result, GPT 4o mini altered its behavior by significantly raising its cooperation rate to 91%. By being overly cautious, GPT 4o mini exposed itself to potential exploitation and subsequently suffered large losses (an average of -4.795 points per round). On the other hand, DeepSeek and Gemini were able to use cooperative and non-cooperative behaviors in combination and maintain relatively low cooperation rates, with 17% and 15.5% respectively. Both models were able to mitigate regulatory costs while still profiting from competitive actions.

5. Conclusions

The empirical results provide important insight into the interactions of pricing algorithms (as related to algorithmic pricing), market behavior, and corporate governance. Moreover, the results establish a clear behavioral framework for autonomous agents competing in oligopoly. Longitudinal simulations also confirm that, if left unchecked in a highly concentrated market environment, pricing algorithms will naturally converge on tacit collusion that maximizes overall profits. In fact, this phenomenon creates significant challenges to traditional strategic management because profit-maximizing directives for corporations can, under certain conditions, lead to serious antitrust violations or other legal issues. On the other hand, adding periodic dynamic regulatory shocks into the simulation demonstrate that strictly monitored and punitive regulatory schemes can easily disrupt these collusions. Disruption of such collusive behavior by regulatory action will force a fundamental review of algorithmic governance. The results of this study clearly show that regulatory action is not simply an exogenous constraint on the operation of autonomous systems, it rather fundamentally alters the strategic options available to firms that participate in such systems.

From a strategic management standpoint, one of the most important contributions of this study relates to revealing the inherent trade-offs between compliance and aggressiveness in autonomous corporate systems. That is, while very aggressive pricing models can produce large increases in revenue in low-over-sight environments, they tend to be structurally rigid and thus extremely vulnerable to catastrophic losses due to regulatory fines at times of regulatory transition. As such, active compliance management must evolve from being a relatively passive corporate cost-center to becoming a critical component of strategic decision-making. Corporate leaders can no longer view the deployment of algorithms as solely an operational issue and thus treat the deployment of algorithms as an entirely decentralized, technical issue. Rather, corporate leaders must internalize the findings of this study and embed algorithmic auditing and legal guard-rails directly into the highest-level risk-management processes of their respective organizations so that their autonomous pricing mechanisms remain legally compliant and ethically responsible in accordance with broader notions of corporate social responsibility.

There are many potential applications of this study to both future organizational design and future regulatory design. For example, firms may use the simulated stochastic model developed here to test their proprietary pricing models prior to actual deployment for potential antitrust violations. Moreover, there are many potential ways to extend this methodology to enable "compliance-by-design" architectures wherein ethical and legal constraints are included in the objective function of the algorithm along-side profit objectives. Future directions for this line of research include developing hybrid corporate forms that explore how collaborative frameworks involving humans can mediate between the need for autonomous market responsiveness and human oversight to ensure that unintended collusion does not occur.

While the current study provides a number of important insights regarding the nature of autonomous systems, it is essential to acknowledge several key limitations of the current research. First, all of the simulations in this study were based upon a stylized oligopolistic competition model that has been constrained to have limited operational dimensions. These simplifications may distort the complex macro-economic factors, idiosyncratic demand shock variability and supply-chain disruptions present in actual markets. Second, it was assumed that regulators would detect violations uniformly and perfectly across time. In reality, antitrust enforcement is often delayed, diffuse and subject to protracted litigation appeals. Thirdly, the behavioral data used in this study is based on a single point-in-time versions of LLMs. It follows that observed strategic responses may relate to specific version(s) of training configuration(s) rather than representing some permanent/universal rule governing the evolution of LLMs. Future studies should seek to address these limitations through development of multi-agent competitive environments with varying levels of detection probability for regulators and inclusion of human decision makers in the competitive loop.

References

1. Affonso, F.M. (2026) ‘Large language models converge on competitive rationality but diverge on cooperation across providers and generations’, *arXiv*. Available at: https://arxiv.org/abs/2604.18596 (Accessed: 13 May 2026).
2. Fish, S., Kraus, S., Malone, T., & others (2024) ‘Algorithmic Collusion by Large Language Models’, *American Economic Association Papers and Proceedings*. Available at: https://www.aeaweb.org/conference/2025/program/paper/GDskRTN3 (Accessed: 13 May 2026).
3. Gandhi, K., Sadigh, D., & Goodman, N.D. (2023) ‘Strategic Reasoning with Language Models’, *arXiv*. Available at: https://arxiv.org/abs/2305.19165 (Accessed: 13 May 2026).
4. Hadfield, G.K., & Clark, J. (2025) ‘An Economy of AI Agents’, *NBER Working Paper*. Available at: https://www.nber.org/system/files/chapters/c15305/revisions/c15305.rev1.pdf (Accessed: 13 May 2026).
5. Hua, Y., Qu, L., & Haf, R. (2024) ‘Assistive Large Language Model Agents for Socially-Aware Negotiation Dialogues’, *Findings of the Association for Computational Linguistics: EMNLP 2024*, pp. 8047-8074. Available at: https://doi.org/10.18653/v1/2024.findings-emnlp.473 (Accessed: 13 May 2026).
6. Idowu, J. (2026) ‘Mapping Human Anti-collusion Mechanisms to Multi-agent Artificial Intelligence Systems’, *Knowledge-Based Systems*. Available at: https://doi.org/10.1016/j.knosys.2026.116067 (Accessed: 13 May 2026).
7. Jackson, M.O., & Zenou, Y. (2015) ‘Games on networks’, *Handbook of Game Theory*, 4, pp. 95–163. Available at: https://www.nber.org/papers/w15984 (Accessed: 13 May 2026).
8. Keppo, J., Li, Y., Tsoukalas, G., & Yuan, N. (2026) ‘On the Fragility of AI Agent Collusion’, *SSRN Electronic Journal*. Available at: https://www.researchgate.net/publication/402573701_On_the_Fragility_of_AI_Agent_Collusion (Accessed: 13 May 2026).
9. Kolter, Z. (2025) ‘The AI Agent Era Requires a New Kind of Game Theory’, *Wired Magazine*. Available at: https://www.wired.com/story/zico-kolter-ai-agents-game-theory (Accessed: 13 May 2026).
10. Krakowski, S., Luger, J.,& Raisch, S. (2023) 'Artificial intelligence and the changing sources of competitive advantage', Strategic Management Journal, 44(6), pp. 1425-1452. Available at: https://doi.org/10.1002/smj.3387 (Accessed: 16 May 2026).
11. Lasry, J.M., & Lions, P.L. (2007) ‘Mean field games’, *Japanese Journal of Mathematics*, 2(1), pp. 229-260. Available at: https://doi.org/10.1007/s11537-007-0657-8 (Accessed: 13 May 2026).
12. Lebovitz, S., Lifshitz-Assaf, H., & Levina, N. (2022) 'To Engage or Not to Engage with AI for Critical Judgments: How Professionals Deal with Opacity When Using AI for Medical Diagnosis', Organization Science, 33(1), pp. 126-148. Available at: https://doi.org/10.1287/orsc.2021.1549 (Accessed: 16 May 2026).
13. Motta, M. (2004) *Competition Policy: Theory and Practice*. Cambridge: Cambridge University Press.
14. Nash, J. (1950) ‘Equilibrium points in n-person games’, *Proceedings of the National Academy of Sciences*, 36(1), pp. 48-49. Available at: https://www.jstor.org/stable/88031 (Accessed: 13 May 2026).

15. Nikandrova, A., & Parekh, A. (2025) ‘Algorithmic Exclusion by Large Language Models’, *SSRN Electronic Journal*. Available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5550761 (Accessed: 13 May 2026).
16. Nisan, N., Roughgarden, T., Tardos, É., & Vazirani, V.V. (eds.) (2007) *Algorithmic Game Theory*. Cambridge: Cambridge University Press. Available at: https://www.cs.cmu.edu/~sandholm/cs15-892F13/algorithmic-game-theory.pdf (Accessed: 13 May 2026).
17. Payne, K., & Alloui-Cros, B. (2025) ‘Strategic Intelligence in Large Language Models: Evidence from Evolutionary Game Theory’, *arXiv*. Available at: https://arxiv.org/abs/2507.02618 (Accessed: 13 May 2026).
18. Sashihara, J., Fujita, Y., Nakamura, K., Terui, A., & Hayashi, T. (2025) ‘LLM-based Multi-Agent System for Simulating Strategic and Goal-Oriented Data Marketplaces’, *arXiv*. Available at: https://doi.org/10.48550/arXiv.2511.13233 (Accessed: 13 May 2026).
19. Suzuki, R., & Arita, T. (2024) ‘Artificial intelligence can evolve into more selfish or cooperative personalities through game theory and large-scale language models’, *Nagoya University Research News*. Available at: https://en.nagoya-u.ac.jp/news/articles/research_information_201/ (Accessed: 13 May 2026).
20. Vallinder, A., & Hughes, E. (2025) ‘Cultural Evolution of Cooperation among LLM Agents’, *Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025)*, pp. 2771–2773. Available at: https://arxiv.org/abs/2412.10270 (Accessed: 13 May 2026).
21. Von Neumann, J., & Morgenstern, O. (1944) *Theory of Games and Economic Behavior*. Princeton: Princeton University Press. Available at: https://archive.org/details/theoryofgameseco00vonn (Accessed: 13 May 2026).
22. Yao, J., Chen, C., & Zhang, B. (2026) ‘Competition and Cooperation of LLM Agents in Games’, *arXiv*. Available at: https://arxiv.org/abs/2604.00487 (Accessed: 13 May 2026).
23. Zhang, Y., Mao, S., Ge, T., Wang, X., de Wynter, A., Xia, Y., & Wei, F. (2024) ‘LLM as a Mastermind: A Survey of Strategic Reasoning with Large Language Models’, *arXiv preprint arXiv:2404.01230*. Available at: https://doi.org/10.48550/arXiv.2404.01230 (Accessed: 13 May 2026).
24. Zhu, K., Du, H., Hong, Z., Yang, X., Guo, S., Wang, Z., & You, J. (2025) ‘MultiAgentBench: Evaluating the Collaboration and Competition of LLM Agents’, *arXiv*. Available at: https://doi.org/10.48550/arXiv.2503.01935 (Accessed: 13 May 2026).

APPENDIX

Appendix A. Supplementary Data

This appendix contains the complete dataset generated during the simulation, including the overall leaderboard, a segmented phase by phase analysis and the chronological logs capturing every strategic decision made across the 1,000 round simulation. The dataset is provided as a supplementary Excel file submitted together with the paper, allowing full transparency and replication of the reported results.

Oligopoly_Final_1000Rounds.xlsx

Data and Code Availability

The datasets generated and analyzed during the current study, together with the full Python code used to implement the simulation and process the results, are publicly available in a GitHub repository. The repository includes all necessary files to reproduce the experimental setup, run the simulation, and verify the findings: https://github.com/ruxi-stnmr/oligopoly-simulation/tree/main