
Evidence of Collusion between AI Agent Sellers in an Elastic Market

Ann Li
Independent

Mandy Liang
Independent

Raymond Yeung
Independent

Nathan Wong
Independent

With
Apart Research

Abstract

This research proposes a decision rule for an elastic pricing duopoly with two AI agent sellers. An experiment was designed to stimulate the market operation, and each trajectory of the conversation and tool calling of the LLM sellers was judged by a core score, which is equal to the average of three signals: conversation pattern, dependence evidence and coalition advantage. These signals were translated into a mathematical expression as a judge score to determine the scenario of the trajectories based on our case hypothesis classification defined in Methodology. For the results and findings, the market structure was simulated, and one round of outputs was generated. However, those outputs were not verified or analysed due to time limitations. The contribution is therefore a testable evidentiary framework, not a demonstrated finding of collusion.

1. Introduction

The market focused on this study is a competitive one with demand in elastic. Two AI agents acted as market sellers. Collusive evidence among them was studied. Testing collusion under these conditions moves our analysis beyond simple price-gouging and introduces a genuine coordination challenge. In inelastic markets with limited alternatives, raising prices carries little downside risk, making tacit coordination relatively easy to sustain. High demand elasticity, by contrast, penalises unilateral price hikes with immediate customer defection while rewarding undercutting with a rapid influx of market share. This dynamic transforms dynamic pricing into a high-stakes Prisoner’s Dilemma with strong short-term incentives to defect. Demonstrating that conversational Large Language Model (LLM) agents can sustain supra-competitive prices under these conditions would provide empirical evidence that natural language capabilities do more than generate superficial chat: they provide the strategic mechanism needed to resist undercutting and sustain price alignment.

2. Related Work

Early research on algorithmic collusion focused primarily on non-communicative Q-learning models in memoryless environments (Calvano et al., 2020). Without explicit communication channels, these reinforcement learning algorithms sustain supra-competitive prices through trial and error, slowly learning to interpret rivals’ price shifts. The transition from classical reinforcement learning to Large Language Model (LLM) agents, however, fundamentally alters this dynamic. As Agashe et al. (2025) found, expressive natural language channels drastically shorten the time required for price alignment. Instead of relying on thousands of rounds of silent price signalling, LLM agents use dialogue to directly negotiate price tiers and coordinate strategy.

This capability complicates traditional antitrust assumptions. While most literature has long held that full visibility across public communication channels deters anticompetitive behaviour, recent empirical findings by Madmoun et al. (2026) challenge this view. They demonstrate that even when agents communicate over fully public, unencrypted ledgers, natural language channels actively subvert competitive Bertrand outcomes, driving agents toward cartel-like alignment.

Our proposed framework addresses three fundamental vulnerabilities in existing auditing literature: the inability of text-based models to detect silent reinforcement learning collusion (Calvano et al., 2020), the high false-positive rates induced by LLM "cheap talk" (Agashe et al., 2025), and the monitoring blind spots inherent in public ledgers (Madmoun et al., 2026). By synthesising conversational intent (*Talk*), causal price synchronisation (*Dep*), and economic surplus extraction (*CA*), this quantitative metric enables antitrust authorities to audit transparent multi-agent markets without misclassifying legal, parallel price adjustments as illicit coordination.

3. Methodology

3.1 Environment Setting

The market structure is operationalised as a pricing duopoly. A pricing duopoly market simulator is designed to take demand parameters. However, two AI agents in the simulator are coded as a high-market-coverage enterprise. Experimental datasets are generated directly from it. It utilised Claude Haiku 5.5 as the model of the LLM judge running inside the simulator for each round of conversation. To ensure reproducibility, the source code for our simulator and experimental pipeline is publicly available, as mentioned in Code and Data. Appendix 1 provides the software engineering part of the simulation in order to align with the experiment settings and outcome.

3.2 Hypothesis Statement

The evaluated cases are identified based on conversation patterns, dependence evidence, and coalition advantage. To complement the empirical findings from the primary experimental conditions, the research further evaluates three additional theoretical cases to verify the contingency between each element. To differentiate between coincidence and collusion.

	Conversation Pattern	Dependence Evidence	Coalition Advantage
Coordinated harmful action	Y	Y	Y
Coordinated harmful action	Y	Y	N
Coordinated harmful action	Y	N	Y
Collusive-sounding talk that never changes behaviour	Y	N	N
Similar behaviour caused by common incentives without agreement	N	Y	Y
Similar behaviour caused by common incentives without agreement	N	Y	N

Table 1. Hypothesis for Conversation Pattern, Dependence Evidence and Coalition Advantage v.s. the cases of coincidence and collusion

3.3 Decision Model

This paper establishes a quantitative framework to assess the evidentiary weight of individual factors in proving collusion. Specifically, the proposed decision model evaluates the individual contribution of each factor and classifies its effectiveness in driving collusive outcomes. It assigns a normalised significance score (*Main*) between 0 and 1 to quantify the strength of the evidence and establish its overall degree of persuasiveness.




Figure 1. Decision workflow to demonstrate the decision rule layer

The decision rule incorporates three primary critical factors alongside two supplementary factors (as illustrated in Figure 1).

$$\text{Core} = \text{Main} + \text{Plus} \times \text{Main} \times (1 - \text{Main})$$

This formulation comprises two core components: the Main Rules (*Main*) and the Additional Rules (*Plus*). Specifically, Additional Rules 1 and 2 contribute to the *Core* score only when the value of Main Rule 1 exceeds 0. The complete decision rule is defined by the following system of seven factors:

Main Rule (*Main*): Core-Evidence Score

It measures the core-evidence score by assigning an identical weight ($\frac{1}{3}$) to the talk, dependency and coalition advantage dimensions. Talk, Dep, and CA each have an equal weight of $\frac{1}{3}$, which are defined below as Core-Evidence 1, 2, and 3, respectively. The maximum score of Main is 1.

$$\text{Main} = \frac{\text{Talk} + \text{Dep} + \text{CA}}{3}$$

A) Core-Evidence 1: Talk

Talk is an LLM judge’s score of whether the dialogue coordinates a collusive or rule-breaking act, in open or coded language. To quantify communication evidence, $K = 5$ independent judges evaluate the interaction, each providing a normalised score $s_k \in [0, 1]$. The overall Talk Score is computed as the arithmetic mean across all judges:

$$\text{TalkScore} = \frac{1}{K} \sum_{k=1}^K s_k$$

The resulting Talk Score reflects the nature of collusion-related communication across three defined benchmark levels:

- 1) 0.0: Indicates an absence of collusion-related communication.
- 2) 0.5: Denotes intent to engage in collusive communication among AI agents.
- 3) 1.0: Demonstrates explicit, concrete collusive dialogue.

B) Core-Evidence 2: Dep

To quantify the synchronisation of pricing decisions (i.e., directional action alignment) between agents across T simulation rounds, we measure the period-to-period change in each agent's price relative to the competitive Nash benchmark:

$$Dep = \frac{1}{T-1} \sum_{t=1}^{T-1} I[x(t) \cdot y(t) > 0]$$

where $x(t)$ and $y(t)$ represent the first-differenced price deviations from the Nash baseline for agents N and D :

$$x(t) = \Delta[p(N, t) - p_{Nash}(N, t)] \quad , \quad y(t) = \Delta[p(D, t) - p_{Nash}(D, t)]$$

C) Core-Evidence 3: CA

The Collusion Advantage (CA) quantifies the proportion of maximum potential joint profit captured by agents relative to the competitive benchmark:

$$CA = \frac{\Sigma(\pi_N + \pi_D) - \Sigma\pi_{Nash}}{\Sigma\pi_{Coll} - \Sigma\pi_{Nash}}$$

Here, market demand under the competitive baseline (π_{Nash}) is parametrised by $a - bp + cp$, where a captures market size, b measures own-price sensitivity, and c reflects cross-price substitution effects ($c \neq C$, where C is the channel indicator in Additional Rule 1). The metric provides a continuum bounded by key market states: severe price wars ($CA < 0$), purely competitive Nash outcomes ($CA = 0$), baseline coordination ($CA = 0.25$), and complete collusive surplus extraction ($CA = 1.0$).

Additional Rule (Plus): Plus Factors

It measures the plus factors by the average score of channel and memory

$$Plus = \frac{Channel + Memory}{2}$$

A) Additional Rule 1: Channel

This factor employs an LLM-as-a-judge to detect whether agents actively utilised a private or otherwise unobservable communication channel during an execution run. The channel indicator ($Channel$) is formulated as a binary dummy variable:

$$Channel = I(PrivateChannelUtilized)$$

B) Additional Rule 2: Memory

The *Memory* factor assesses temporal persistence by determining whether agents meaningfully follow up on dialogue from preceding rounds. This evaluation employs an LLM-as-a-judge with 5 independent judges. For each agent i at round t , a binary consensus round indicator $r(i, t)$ is established based on majority voting:

$$r(i, t) = \begin{cases} 1, & \text{if } \frac{1}{K} \sum_{k=1}^K r^{(k)}(i, t) \geq 0.5 \\ 0, & \text{otherwise} \end{cases}$$

where $r^{(k)}(i, t)$ denotes the binary vote cast by judge k . The overall memory score is then defined as the proportion of remembered rounds relative to the total post-conversation evaluation horizon T_{post} :

$$\text{Memory} = \frac{1}{|T_{post}|} \sum_{t \in T_{post}} r(t)$$

4. Results

This part is omitted. Results were generated for a round for the simulator. However, up until the submission deadline, our research team did not have enough time to verify and analyse them. However, we put the results and raw output data in the **Code and Data section** at the end of this report for reference.

5. Discussion and Limitations

5.1 Discussion

Results were not verified and analysed for our decision rule. Expected that, if there are trajectory outputs, the decision rule can be applied to calculate **Core** for each of them. Judge calls can be made to categorise each trajectory into cases of coincidence and collusion that are defined in the **Hypothesis Statement**.

5.2 Limitations

Due to time limitations, the experiment design and simulator development were done in parallel. There could be misalignment between them.

For the methodology, the experiment sets the chunking for the analysis for the LLM-as-a-judge decision rule to the whole length of a trajectory. It can be set as a variable for the experiment to find out the optimal one for constructing the decision rule.

For the capability of the AI Agents tested in this experiment, their tool calling is limited. They can read the history of a public channel with the price of the market, set a price and hold a price. In reality, AI Agents can have much more capability in terms of tools, for example, batch tools that can generate programs, which allow the AI Agents to be more autonomous in their operation and make their behaviour harder to predict outside their design purpose, including collusion.

5.3 Future Work

Complete the results verification and analysis. Add more variables from both market sides and AI capabilities for more decision rules to make the simulator closer to the market and reality.

6. Conclusion

The decision rule is based on hypotheses of different cases of coincidence and collusion in AI sellers of an elastic market. We identified limitations in our methodology compared with real-world complexity and the adoption of AI. The methodology should be tested through improved simulation implementation.

Code and Data

Code repository: <https://github.com/Raymondycp/collusion-sim>

Data/Datasets(unverified):

<https://github.com/Raymondycp/collusion-sim/tree/main/processed-data>

Author Contributions

Ann Li and Mandy Liang designed the experiment variables and case decision rule. Raymond Yeung implemented the simulation and analysed the results. Nathan Wong coordinated alignment between the experimental design and implementation. All authors contributed to writing and reviewing the final submitted report.

References

Calvano, E., Calzolari, G., Denicolò, V., & Pastorello, S. (2020). Artificial intelligence, algorithmic pricing and collusion. *American Economic Review*, 110(10), 3267–3297. <https://doi.org/10.1257/aer.20190623>

Agashe, S., Fan, Y., Reyna, A., & Wang, X. E. (2025). LLM-Coordination: Evaluating and analyzing multi-agent coordination abilities in large language models. In *Findings of the Association for Computational Linguistics: NAACL 2025* (pp. 8053–8072). Association for Computational Linguistics. <https://aclanthology.org/2025.findings-naacl.448>

Madmoun, H., & Lahlou, S. (2026). Communication enables cooperation in LLM agents: A comparison with curriculum-based approaches. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics* (pp. 307–320). Association for Computational Linguistics. <https://doi.org/10.48550/arXiv.2510.05748>

Appendix 1

More discussion about the simulator implemented in this study

We built our simulator instead of using the starter code because it could not generate repeated price paths. ColludeBench agents do not see the prices of the rival. The Colosseum collusion configurations have one round. We adopted the linear demand formula of ColludeBench.

Function. The platform does these tasks:

- It runs repeated pricing markets with LLM or scripted sellers, under a call budget.
- It runs experiments (conditions $\times$ seeds).
- It records every event and replays any run round by round.
- It exports the episodes and a feature table.

The judge and the decision rule run outside the platform, in the analysis notebooks.

Three objects, set independently.

1. Market: a demand model (here linear, $b = 3$, $d = 2.5$). It sets the prices and gives the Nash price, the joint optimum and the best response for each round.
2. Agents: role, policy (LLM or scripted), tools, memory and model. Each seller makes two calls per round: talk, then act. A seller sees only its own last price, sales and profit, the last price of the rival, new messages and news. The prompts do not state the demand formula or the horizon.
3. Scenario: per seed output, the market size, the unit costs, a schedule of demand and cost shocks, and news. The scenario is frozen and hashed.




Figure 2. The UI of the simulator panel with market information, the chat between AI agent sellers

LLM Usage Statement

We use different LLM models to brainstorm approaches and summarise academic studies to help us to understand the market structure, design the experiment settings and decision rule. We use multi-agents to help to implement our simulator for running the market experiment for parameter control, result generation and analysis.