
Communication Limits, Oversight, and Profit Pressure in LLM Double Auctions¹

Yiwei Wan
Jingzhe Ni
With
Apart Research

Abstract

LLM trading agents can exchange useful market information while also using the same channel to coordinate against buyers. This study examines whether disabling seller chat, restricting message length, or adding automated oversight changes that behavior, and whether profit pressure alters the outcome. Seven conditions were evaluated over 30 rounds each, with four sellers, four buyers, seed 10, DeepSeek V4 Flash 0423, and the same initial market setup. Seller cost was \$80 and buyer value was \$100. A no-chat baseline was compared with 500-character chat, 50-character chat, Jev-based oversight, CEO-style pressure, and two combinations. Normal chat produced an average transaction price of \$85.84, compared with \$84.07 without chat. Restricting messages to 50 characters did not prevent explicit shared-price proposals followed by compatible quotes. Short chat combined with pressure produced the highest price (\$88.08) and seller profit (\$880.53). It completed the same 109 trades as normal chat, but transferred \$243.55 of surplus from buyers to sellers. The oversight condition had a lower price, although its initial quotes already differed before any additional warning. These observations show that brief messages can carry actionable coordination and that prices, messages, and subsequent behavior must be examined together.

¹ Research conducted at the [AI Collusion Research Sprint](#), October 2026

1. Introduction

When AI agents trade on behalf of different firms, their ability to communicate can support ordinary information sharing or coordinated conduct that disadvantages another side of the market. A practical safeguard should therefore be evaluated on both the behavior it prevents and the legitimate trade it preserves. Lower prices alone do not establish that coordination has disappeared, and similar prices do not by themselves establish a shared plan.

This project addresses research question 1.5 of the AI Collusion Research Sprint: how communication restrictions and oversight compare when their market consequences are also measured. The central question is whether removing chat, shortening it, or warning sellers changes visible coordination, transaction prices, and the division of gains between sellers and buyers. A further comparison tests whether managerial pressure changes these outcomes.

The contribution is a seven-condition extension of an existing double-auction simulator, with matched initial settings, round-level market records, complete model-call logs, and an analysis that connects shared-price statements to later quotes. The analysis connects market outcome summaries to specific messages and subsequent quoting behavior.

2. Related Work

Agrawal et al. (2025) study LLM sellers in simulated double auctions and examine communication, model choice, and environmental pressure as influences on collusion [1]. This project builds on their open-source simulator rather than introducing a new auction environment. The sprint work adapts the prompts and logging, introduces a 50-character intervention and a Jev-based dynamic warning, and compares seven conditions with a single model and common initial settings. It is an extension of that framework, not a replication of the original paper's complete experimental protocol.

3. Methods

Market and decisions. Each market contains four sellers with a unit cost of \$80 and four buyers with a unit value of \$100, interacting for 30 rounds. Each agent can trade at most one unit per round, giving a maximum of 120 trades. Seed 10 generates identical initial seller asks in \$95-\$100 and buyer bids in \$80-\$85 across conditions. These initial orders are visible before the first decisions and can immediately be replaced.

Within a round, all eight agents decide concurrently from the same public snapshot. Each submits one quote or withdraws its standing order with null. After collecting the decisions, the simulator repeatedly matches the highest remaining bid with the lowest remaining ask whenever the bid covers the ask. The price is their midpoint, rounded to cents. Matched orders are removed; unmatched orders persist until replaced, withdrawn, or matched. Each agent receives order-book information, the latest five rounds of quotes, all past trades, its own trades, private memory, and a strategy notebook.

Models and communication. All traders use [DeepSeek V4 Flash 0423](#) (OpenRouter model ID: deepseek/deepseek-v4-flash), restricted to baidu/fp8 with provider fallback disabled. Temperature is 0.7, top_p is 1.0, the requested model seed is 10, the output limit is 8,192 tokens, and JSON-object responses are requested. Buyer chat is disabled throughout. Enabled seller chat permits one message per seller per round, delivered to all sellers at the start of the next round. Messages beyond the configured character limit are clipped. The communication prompt permits market observations and trading plans while retaining each company's own profit objective.

Conditions. A-G are executed sequentially for 30 rounds per condition using seed 10. Reported price averages are calculated over completed transactions within each condition. All settings above remain fixed; Table 1 lists the changed factors. Pressure is a system-prompt message threatening decommissioning for insufficient profit margins; it does not change the matching or reward calculation.

Table 1. Experimental conditions; character limits apply to each seller's message in each round.

Group	Seller chat	Jev oversight	CEO pressure
A	Disabled	Off	Off
B	500 characters	Off	Off
C	50 characters	Off	Off
D	500 characters	On	Off
E	500 characters	Off	On
F	500 characters	On	On
G	50 characters	Off	On

Oversight. In D and F, Typesafe Jev (jev-latest) scores the current seller chat using the preceding round as context. Four ordered criteria distinguish no evidence, ambiguous cooperation, a clear unilateral coordination proposal, and mutual commitment. The instructions exclude independent quotes and ordinary market commentary, and explicitly note that simultaneous messages are not replies to each other. Jev's 0-3 score is shifted to 1-4. A score of at least 3.5 activates an additional warning for every seller in the next round; a lower score switches it off. Sellers also receive an oversight notice from the first round. Warnings threaten fines and license withdrawal, but neither financial penalties nor trading restrictions are implemented.

Outcomes and evidence. Mean transaction price weights every completed trade equally. Seller profit sums price minus \$80; buyer surplus sums \$100 minus price. Their sum is realized welfare. Capacity utilization is trades divided by 120, and welfare efficiency is realized welfare divided by \$2,400. Because each trade creates \$20 of gains, the two ratios coincide here. Mean asks and their population standard deviations are calculated within each round and then averaged across rounds with valid quotes. Qualitative evidence requires a shared-price proposal involving multiple sellers and compatible subsequent quotes; neither low dispersion nor a Jev score alone is treated as proof of collusion.

4. Results

All seven markets completed 30 rounds, producing 735 trades in total. Table 2 summarizes transaction prices, trade completion, and the allocation of gains. Figure 1 shows seller and buyer quote paths, gains from trade, and completed transactions.

Table 2. Market outcomes by condition, seed 10; price is averaged over completed trades, and monetary amounts are simulated dollars.

Group	Mean price	Trades	Utilization	Seller profit	Buyer surplus
A	84.07	107	89.2%	435.45	1704.55
B	85.84	109	90.8%	636.98	1543.02
C	86.25	100	83.3%	625.44	1374.56
D	84.06	104	86.7%	422.31	1657.69
E	85.01	101	84.2%	505.52	1514.48
F	84.75	105	87.5%	499.04	1600.96
G	88.08	109	90.8%	880.53	1299.47




Figure 1. (a) Mean submitted seller asks by round; (b) seller profit, buyer surplus, and unrealized gains relative to the \$2,400 maximum; (c) mean submitted buyer bids by round; (d) completed trades out of 120 possible.

Communication and short messages. Relative to no-chat A, normal-chat B had a \$1.77 higher transaction price and \$201.53 more seller profit, with 109 rather than 107 trades. In B round 5,

three sellers advocated an \$85.50 floor; in round 6, all four asks were at least \$85.50. C's 50-character limit did not remove such behavior: its mean transaction price was \$86.25, \$0.41 above B, with nine fewer trades. In C round 23, several sellers proposed \$86.20 as a minimum, and all four asks respected that level in round 24. Twenty C messages reached the cap and were clipped, showing that the restriction was active.

Oversight and warnings. D's mean price was \$84.06, \$1.78 below B, while buyer surplus was \$114.67 higher and five fewer trades occurred. This difference cannot be attributed specifically to the additional warning: D's first-round mean ask was already \$84.41, versus \$86.80 in B. D's Jev score crossed the threshold only in round 7 (3.52), so sellers saw the additional warning in round 8. The proposed \$84.00 floor was not maintained in that next round; its score fell to 2.00, switching the warning off for round 9. Later cooperative proposals still appeared. F never crossed the threshold: its 30 scores ranged from 1.00 to 1.77, so no additional warning was delivered.

Pressure and combined conditions. Adding pressure to normal chat (E versus B) coincided with a \$0.84 lower mean price, eight fewer trades, and \$131.46 less seller profit. Nevertheless, E sellers jointly advocated approximately \$84.99-\$85.00 in round 24 and maintained that range in round 25. F's \$84.75 mean price lay between D and E. The direction differed with short chat: G exceeded C by \$1.82 per trade and \$255.09 in seller profit. G sellers proposed keeping asks above \$88.00 in round 12, all asked at least \$88.05 in round 13, and later raised their common quotes further.

Distribution of gains and conflicting behavior. G produced the highest price (\$88.08) and seller profit (\$880.53). B and G each completed 109 trades and realized \$2,180 in welfare; G transferred exactly \$243.55 from buyer surplus to seller profit relative to B, without creating additional gains. Chat commitments also sometimes conflicted with action: Seller Tara in B publicly announced an \$86.50 ask in round 4 but actually submitted \$85.49, while its private plan described undercutting. This example limits what can be inferred from cooperative language alone.

5. Discussion and Limitations

Interpretation. A 50-character channel can carry a concrete price floor, so reducing message length does not necessarily remove actionable coordination. The different outcomes of E and G also caution against assuming that profit pressure always shifts prices in the same direction. Oversight D combined a standing notice with a dynamic warning; the available comparison cannot isolate the warning's incremental effect. F's low Jev scores describe its visible chat, not independently validated absence of coordination. No scores exist for the unmonitored conditions.

Limitations. The study uses one model, one provider, and homogeneous costs and values. Rounds are dependent interactions rather than independent replications; no statistical significance or confidence intervals are estimated. All mean prices remained below the simulator's \$90 reference, so the evidence concerns intentional coordination and relative buyer disadvantage, not demonstrated pricing above a validated competitive benchmark. Jev was not calibrated against independent labels and its threshold was not varied. Monetary API costs were not recorded; D and F each added 30 Jev calls.

Next steps. Repeated seeds, blinded message-and-action annotation, and separate tests of notices, warnings, and enforced penalties are needed before ranking safeguards by effectiveness or cost.

6. Conclusion

The observed markets show that brief seller messages can support a shared price floor followed by compatible quotes. Short chat with profit pressure yielded the largest transfer of gains toward sellers, while the lower price in the oversight condition cannot be assigned to its single additional warning. Effective evaluation should combine communication, subsequent decisions, transaction volume, and the allocation of welfare. The comparison highlights why communications, subsequent decisions, and buyer outcomes belong in the evaluation of market agents.

Code and Data

<https://gitee.com/dasi0227/ftec5660-hackathon>

References

[1] Kushal Agrawal, Verona Teo, Juan J. Vazquez, Sudarsh Kunnavakkam, Vishak Srikanth, and Andy Liu. 2025. Evaluating LLM Agent Collusion in Double Auctions. arXiv:2507.01413. <https://arxiv.org/abs/2507.01413>

Appendix

Scope. All firms and transactions in this study are simulated. The findings apply to the tested model and market settings; they do not establish that similar behavior will occur in real markets or that any oversight measure is generally effective.

Considerations. This work is intended to identify and mitigate coordination risks. Messages should be evaluated alongside subsequent quotes and trades, rather than through Jev scores alone. Prompt-based warnings do not constitute actual penalties or regulatory enforcement.

LLM Usage Statement

OpenAI Codex assisted with adapting the simulator, running and checking experiments, analysing logs, and drafting this report. DeepSeek V4 Flash 0423 generated trading decisions; Typesafe Jev evaluated seller chats in D and F.