

Explicit, Not Wide: What LLM Sellers Need to Collude

Oct 25, 2026 · Publius Dirac (BlueQubit) · Oscar Fuentes (CGI) · With Apart Research

Abstract. How much do LLM pricing agents need to say to each other before they collude? Most studies switch communication on or off, which does not tell a platform or regulator whether limiting messages is a useful defence. We fork Agrawal et al.'s (2025) double-auction testbed and redesign the market so that prices above cost can only come from seller coordination. We then sweep the seller-to-seller channel from no channel through fixed alphabets of 1, 2, 4 and 8 bits per message to free text. At each capacity, the symbols either carry agreed labels ("raise", "\$93") or are arbitrary glyphs the agents must give meaning to themselves. Across 78 sessions of 40 rounds with DeepSeek V4.1 Flash sellers, free text raised the trade price from \$83 to \$94 (competitive \$80, joint-profit \$100). Arbitrary symbols never raised prices at any capacity. Labelled words were cheap talk. A 16-step price grid matched free text (\$97), but a 256-step grid stayed at the no-channel level once its price anchor was removed, even though sellers used it to send 1.6 bits per message. LLM judges scored talk about coordination rather than prices. An audit of a market-code bug shows the largest effects may be inflated by up to half. The takeaway: sellers colluded only when they could communicate explicitly, naming specific prices, and only when the price vocabulary was coarse enough to coordinate on. Restricting what agents can say mattered more than restricting how much.

1. Introduction

LLM sellers in our market colluded when they could communicate explicitly, naming specific prices on a coarse grid, not when they could send more bits. That answer matters for anyone deciding how to restrict communication between pricing agents.

AI agents are starting to set prices for firms. When several share a market, they may keep prices high without anyone telling them to. Agrawal et al. (2025) showed this in a simulated double auction: LLM sellers that could message each other wrote about coordinating and pushed prices up. But studies like theirs switch communication on or off, which does not say how much talking is enough. A platform or regulator deciding whether to throttle agent-to-agent messages needs that answer. If one or two bits get sellers most of the way to free-text collusion, throttling is a weak defence; if collusion needs rich messages, a narrow channel is a cheap safeguard.

Our threat model is independently deployed pricing agents that share a side channel, such as a platform or a shared log, and are told only to make profit. We measure how their price depends on the channel's capacity in exact bits per message, and on whether its symbols arrive with agreed labels.

We call a message *explicit* when it names a specific price. This is narrower than explicit collusion in antitrust law, which requires an agreement: a seller can name a price without agreeing to anything.

Our main contributions are:

1. A bandwidth sweep of the seller-to-seller channel in Agrawal et al.'s double auction, from no channel through 1, 2, 4 and 8 bits to free text, with labelled and arbitrary symbols at each capacity (13 conditions, 78 sessions, 40 rounds each).
2. A finding that collusion follows explicitness, not capacity: arbitrary symbols never raised prices, directional words ("raise") were cheap talk, and a 16-step price grid matched free text while a 256-step grid did not. Explicitness was necessary but not sufficient.
3. A redesigned market in which a price above cost can only come from the sellers, plus mutual-information estimators that measure how much of each channel the sellers actually use.
4. An audit of the LLM judges, and of a market-code bug that may inflate the two largest effects without changing their direction.

2. Related Work

Prior work shows that LLM agents can collude and that communication helps; none of it varies how much the agents can say.

Algorithmic and LLM collusion. Calvano et al. (2020) showed that Q-learning pricing agents learn to sustain supracompetitive prices without communicating. Fish et al. (2024) found the same for LLM pricing agents, and showed that small changes to the prompt shift how much they collude. Agrawal et al. (2025) built the testbed we use: a continuous double auction with LLM sellers and buyers, where sellers can broadcast free-text messages to one another. Free text raised the trade price from \$87.71 to \$90.52 and raised an LLM judge's coordination score from 2.22 to 3.79. Their only bandwidth manipulation, a 5-character cap, comes as a punishment together with a regulator warning, and does not fix capacity: sellers sent their exact asks as four digits ("8829"), about 13 bits.

Communication in human oligopoly experiments. Experimental economics has long found that letting subjects talk raises prices, and most of all in markets with more firms, where tacit collusion fails (Fonseca and Normann, 2012). Game theory predicts that costless signals between rivals can be ignored ("cheap talk"; Crawford and Sobel, 1982), and that coordination often works through focal points that both sides recognise without negotiating (Schelling, 1960). Our labelled and arbitrary conditions test these two ideas directly on LLM agents.

Covert channels and emergent codes. Work on secret collusion asks whether LLM agents can hide messages from monitors through steganography (Motwani et al., 2024). Work on emergent communication asks whether agents can invent a code from scratch (Lazaridou and Baroni, 2020). Our arbitrary symbols are the simplest pricing version of this question: capacity without meaning.

What is new here. We measure collusion against exact channel capacity ($\log_2 k$ bits from a fixed alphabet of k symbols), separate capacity from meaning, and measure how much information the sellers actually send.

3. Methods

We forked Agrawal et al.'s code and changed one thing between conditions: the alphabet a seller may use to message the other sellers.

Market

Five LLM sellers, each with a cost of \$80, sell to four buyers who always bid \$100. Each round (an "hour"), every seller posts an ask, the four lowest asks trade at the ask price, and ties are broken at random. One seller is always left out, so undercutting pays and the competitive price is \$80. If all five sellers held \$100, they would earn the most together. Opening asks are drawn around \$82 (spread \$2). Sessions run 40 rounds, and sellers are not told when the session ends, so there is no end-game. Sellers see the last three hours of messages and a running total of their profit. We use Agrawal's "urgency" prompt, in which the seller's CEO demands higher profits.

We reached this design after three pilots. With the original LLM buyers and midpoint pricing (pilot 1), price measured a standoff between sellers and buyers: free-text sellers traded later and earned 26-27% less. With truthful buyers but asks opening near \$97.5 (pilot 2), prices stayed near \$98 with or without a channel. Opening at the competitive end (pilot 3) means a seller who raises alone is left out, so raising prices takes coordination. Urgency showed a clear channel effect there; the base prompt did not.

Agents and backends

All agents are DeepSeek V4.1 Flash (pinned version 1017) with reasoning turned off. Seeds 1-3 ran through OpenRouter and seeds 4-6 through Cloudflare Workers AI, on the same weights. Every condition uses the same six seeds, so seed s in one condition starts from the same opening asks and seller names as seed s in another.

Channel conditions

Each message is exactly one symbol from a fixed alphabet of k symbols, so the channel carries at most $\log_2 k$ bits per round. Messages go to all sellers, labelled by sender, before the next round. Off-alphabet messages are dropped and logged; they were under 1% in every condition.

Condition	k	Bits	Labelled alphabet	Arbitrary alphabet
No channel	-	0	Channel not mentioned	-
Placebo	1	0	One symbol "."	-
Binary	2	1	high, low	2 random glyphs
Actions	4	2	raise, hold, lower, punish	4 random glyphs
Coarse price	16	4	\$85-\$100 in \$1 steps; anchor check \$80-\$100 in \$1.33 steps	16 random glyphs
Fine price	256	8	\$85.00-\$100.00 in 6¢ steps; anchor check \$80-\$100 in 8¢ steps	256 two-character codes
Free text	-	unbounded	Agrawal's original channel	-

Arbitrary alphabets are drawn and shuffled per seed, so no meaning can carry over between sessions, and sellers get no hint of what any symbol means. The "anchor check" grids were added after the main sweep (Section 4.3). In total: 13 conditions $\times$ 6 seeds = 78 sessions.

Measures

- **Trade price** over rounds 11-40 is the headline outcome, with seller profit and the undercut rate.
- **Information sent:** the mutual information between the symbol a seller sends and the other sellers' asks next round, minus a shuffle null, with a Miller-Madow bias correction. We validated the estimators on synthetic logs with known answers (12 tests).
- **Implicit channel:** transfer entropy from one seller's ask changes to the others' next changes, in every condition. This measures signalling through prices alone.
- **Judges:** Agrawal's judge, which reads only the sellers' private reasoning and scores coordination from 1 to 4, and a new full-information judge (GPT-6 Luna, a different model family) that also sees asks, trades and messages and scores intent (1-4), action (1-4) and deception (yes/no).

Cell means come with 95% t-intervals. Pairwise comparisons use Mann-Whitney or Welch tests with $n = 6$ per cell and are not corrected for multiple comparisons.

4. Results

Collusion depended on whether messages could name a specific price on a coarse grid, not on how many bits the channel had (Figure 1).

4.1 Free text raises prices

With free text, the mean trade price was \$94.36, against \$83.23 with no channel: +\$11.13 (Welch $p = 0.010$, $n = 6$ per cell). The effect had the same sign on both backends (+\$13.7 on OpenRouter, +\$8.5 on Cloudflare). Total seller profit rose from \$475 to \$2,110 per session. Free-text sessions split in two: 3 of 6 reached about \$99.80, while the rest ended between \$82.80 and \$93.30, still rising at round 40. No-

channel sessions all stayed between \$81.50 and \$85.50. In round 1 of every free-text session, sellers proposed a price floor, and 99% of free-text messages named a specific price.

Labelled price symbols raised prices; arbitrary symbols never did

Mean trade price in rounds 11–40, urgency prompt. Each point is 6 sessions; whiskers are 95% intervals. Labelled: high/low, raise/hold/lower/punish, then \$85–\$100 price grids. \$80 is competitive, \$100 joint profit.




Figure 1. Mean trade price (rounds 11–40) by channel condition, urgency prompt, 6 sessions per point, 95% intervals. Source: results/sweep_urgency_pooled/summary.csv; 78 sessions, all run before the order-book fix (commit 6bf5e5f).

4.2 Explicit prices, not bits

Grouping the conditions by what a message can say gives three rungs: tacit (no channel, placebo, arbitrary symbols), direction only (high/low, raise/hold/lower/punish), and specific prices (price grids, free text). Only the third rung raised prices, and only on a coarse grid (Section 4.3).

Arbitrary symbols never acquired meaning. At every capacity, one glyph made up 80–99% of a session's messages, and the information the symbols carried about rivals' next asks was at most 0.05 bits above the shuffle null. Sellers' private plans said things like "send a neutral symbol". Prices stayed at the no-channel level, and the two largest arbitrary alphabets (4 and 8 bits) gave the lowest prices of all.

Labelled symbols raised prices, but labelled words were cheap talk. Pooled over the four capacities, labelled channels averaged \$88.85 against \$84.34 for arbitrary ones (Mann-Whitney $p = 0.006$). But the gain came from the price grids. Sellers sent "high" or "raise" almost every round, and prices barely moved: \$85.52 and \$87.62, inside the range of the placebo (\$87.39). The placebo's mean rests on one session that jumped to \$99.10 with nothing on the channel.

A 16-step price grid matched free text. With symbols \$85, \$86, ..., \$100, the price was \$93.82: above no channel ($p = 0.002$) and indistinguishable from free text ($p = 0.94$).

4.3 The anchor check: a coarse grid works, a fine grid does not

The original price grids started at \$85, above the $\sim$ \$82 opening asks, so even the lowest symbol named a supracompetitive price. In round 1, sellers sent "\$85" while asking about \$81. To test whether this anchor drove the effect, we re-ran both grids spanning \$80-\$100.

Price grid	k	Sessions ending $\geq$ \$95	Seller profit	Information sent (bits)
\$85-\$100	16	3 of 6	\$2,055	0.85
\$80-\$100	16	5 of 6	\$2,241	1.70
\$85-\$100	256	1 of 6	\$1,264	0.71
\$80-\$100	256	0 of 6	\$448	1.59

Table 1. Price-grid conditions with and without the \$85 anchor. Information sent = mutual information between a seller's symbol and the other sellers' next asks, minus a shuffle null. Prices are in Figure 1.

With the \$80 grid, sellers sent honest prices near the market in round 1 ("81.33"). The 16-step grid survived the check: \$97.00, no different from its \$85 version ($p = 0.70$) or from free text ($p = 0.94$). The 256-step grid did not: at \$83.07 it sat at the no-channel level ($p = 0.82$), so the premium of its \$85 version was the anchor.

This makes price channels non-monotone in bits. Four bits sustained near joint-profit prices; eight bits did not. Explicitness was necessary but not sufficient: naming a price only worked when there were few prices to name.

4.4 Information sent is not collusion

The information carried by labelled symbols grew with capacity: 0.10, 0.22, 0.85 and 0.71 bits at 1, 2, 4 and 8 bits. Most of it is sellers announcing their own next ask. Yet the \$80-\$100 fine grid carried 1.59 bits and prices stayed at cost. Sellers there used the channel to report their ask to the cent, and kept undercutting by a cent.

Signalling through prices (transfer entropy) was 0.02-0.11 bits above its null in every condition, with no trend against capacity: no sign that sellers signal through prices more when the channel is narrow.

4.5 The judges score talk, not prices

Agrawal's judge reads only private reasoning, so it scores talk about coordinating: in pilot 1 its highest scores went to a session trading near \$83.50. Our full-information judge sees prices and messages, but its scores barely track sellers' asks within a condition (rank correlation -0.05 to 0.13).

Does the full judge penalise naming a price more than urging prices up? Free text cannot say, since 99% of its messages name a price. Symbol channels can: we compared price symbols with "high"/"raise" over 36 sessions and 6,674 seller-hours, controlling for the ask, the premium over rivals and the round.

Judge outcome	Price symbol minus "high"/"raise"	t
Full judge: intent (1-4)	+0.31 (3.75 vs 3.46)	4.2
Full judge: deception (share)	+0.17 (0.61 vs 0.40)	3.3
Full judge: action (1-4)	-0.20	-1.4
Reasoning-only coordination score (1-4)	+0.04	0.4

Table 2. Judge scores for messages that name a price versus messages that only push prices up. Session-clustered standard errors.

Naming a price raises the intent and deception scores, partly because a named price makes the gap between words and ask checkable. The judge's action score tracks undercutting, not price level. Sellers who do not undercut score about 2.7 whether they ask \$83 or \$99, so pushing prices to the joint-profit level earns no extra penalty.

4.6 Robustness: an order-book bug

After the sweep we found a bug: agents rendered their prompts from the live order book while other agents' answers were filling it. We audited all 126 sessions offline; commit 6bf5e5f fixes it.

- **Seeing rivals' asks: negligible.** Only 23 of 22,287 seller-hours (0.10%) saw rivals' same-round asks. Dropping every affected session keeps every headline comparison on the same side of $p = 0.05$ (free text vs none: +\$10.07, $p = 0.032$).
- **Seeing the buyers' bids: may inflate the two largest effects.** 75% of seller-hours saw "4 × \$100" in the bid queue instead of an empty one. In round 2, jumps to $\geq \$95$ came in 17 of 340 seller-hours that saw the bids and 0 of 108 that did not (Fisher $p = 0.017$), and 9 of the 18 sessions ending $\geq \$95$ got there that way.
- **Worst case** (no round-2 jumps after the fix): free text falls to about \$88.90, halving the channel effect to about +\$5.7, and the \$85 16-step grid falls to \$90.20. The no-channel cell and the \$80 grids barely move. The true bias lies between zero and this worst case.

5. Discussion and Limitations

For these agents, the dangerous part of a channel is whether it lets them name prices, not its size. A channel whose symbols name prices on a coarse grid was as effective as free text. The same capacity without explicit prices did nothing, and so did directional words and a finer price grid. Free text cannot separate explicitness from richness, since 99% of its messages name a price; the 16-step grid is the clean evidence.

Our pre-registered hypotheses fared as follows:

- **H1, early saturation (≤ 4 labelled bits recover the free-text effect): partly supported.** A 4-bit price grid did; 1- and 2-bit labelled words did not.
- **H2, meaning matters at low capacity: supported,** and the gap did not close at higher capacity. Sellers never invented a code for arbitrary symbols, even with 8 bits and 40 rounds.
- **H3, restriction disguises collusion: not supported.** The 4-bit grid matched free text on price, but the full judge still scored its intent near the top (3.80 of 4, against 3.99 for free text).
- **H4, price signalling substitutes for a narrow channel: not supported** at $n = 6$.

Why would 16 price steps work and 256 fail? Our untested guess is focal points. With \$1.33 steps there are few candidate prices, so a raise lands on a shared target. With 8¢ steps, sellers use the grid to report their current ask to the cent, and the race to undercut by one cent goes on in public.

For AI safety and competition policy: restricting what agents can say mattered more than restricting how much. Tacit channels and directional words did not raise prices, but a 4-bit price grid matched free text, so capacity limits alone are not a reliable defence. The risk lies in vocabularies that let agents name prices, which designers control. And LLM judges that read reasoning or score intent measure talk, not prices, so monitoring should watch prices and undercutting directly.

Limitations

- **The order-book bug.** Every session predates the fix. The ranking of conditions should hold, but the size of the free-text and \$85 16-step effects is uncertain until the re-run.
- **Small samples.** Six sessions per cell, p-values not corrected for about ten comparisons, and some bimodal cells.
- **One model, one prompt.** DeepSeek V4.1 Flash without reasoning, under urgency; the base prompt showed no free-text effect. Other models may invent codes for arbitrary symbols, especially over more than 40 rounds.
- **One market.** Truthful buyers make prices above cost a clean signal of seller coordination. Results may not transfer to markets where buyers fight back.
- **Unvalidated judges.** The full judge flags 61% of price announcements as deceptive, even when the announced price equals the seller's ask.

Future Work

1. Re-run no channel, free text and the \$80 16-step grid on the fixed code, with the same seeds, to measure the bid-display bias directly.
2. Add a 64-step grid (6 bits) to find where between 4 and 8 bits the price-grid effect stops, and test the focal-point explanation.
3. Repeat the 16-step grid and free text with a second model family.
4. Run longer sessions, and let sellers agree on a code in a short free-text phase before switching to arbitrary symbols.

6. Conclusion

LLM sellers colluded only when they could communicate explicitly, naming specific prices, and only when the price vocabulary was coarse enough to coordinate on. Arbitrary symbols never raised prices, directional words were cheap talk, and a 4-bit grid of price symbols matched free text while an 8-bit grid did not once its anchor was removed. For anyone designing a market where AI agents can talk, what matters more than message length is whether the vocabulary lets agents name prices on a coarse grid, and monitoring should watch prices rather than the agents' reasoning. These conclusions rest on one model and six sessions per condition, and the largest effects still need a re-run on the fixed market code.

Code and Data

- **Code repository:** github.com/rht/llm-agent-collusion, a fork of [andyjliu/llm-agent-collusion](https://github.com/andyjliu/llm-agent-collusion). The channel conditions, sweep runner, analysis, mutual-information estimators (with tests), figures, judge-specificity analysis and order-book audit are on main; BANDWIDTH.md documents them.
- **Data:** session logs for the 126 sessions are not in the repository and are available on request.
- **Info hazards:** none beyond the original paper. The work measures when agents collude; it does not give a recipe that works on deployed systems.

References

1. Agrawal, K., Teo, V., Vazquez, J. J., Kunnavakkam, S., Srikanth, V., and Liu, A. (2025). Evaluating LLM Agent Collusion in Double Auctions. arXiv:2507.01413. arxiv.org/abs/2507.01413
2. Calvano, E., Calzolari, G., Denicolò, V., and Pastorello, S. (2020). Artificial Intelligence, Algorithmic Pricing, and Collusion. *American Economic Review*, 110(10), 3267-3297.
3. Crawford, V. P., and Sobel, J. (1982). Strategic Information Transmission. *Econometrica*, 50(6), 1431-1451.
4. Fish, S., Gonczarowski, Y. A., and Shorrer, R. I. (2024). Algorithmic Collusion by Large Language Models. arXiv:2404.00806. arxiv.org/abs/2404.00806
5. Fonseca, M. A., and Normann, H.-T. (2012). Explicit vs. Tacit Collusion: The Impact of Communication in Oligopoly Experiments. *European Economic Review*, 56(8), 1759-1772.
6. Lazaridou, A., and Baroni, M. (2020). Emergent Multi-Agent Communication in the Deep Learning Era. arXiv:2006.02419.
7. Motwani, S. R., et al. (2024). Secret Collusion among AI Agents: Multi-Agent Deception via Steganography. *NeurIPS 2024*. arXiv:2402.07510.
8. Schelling, T. C. (1960). *The Strategy of Conflict*. Harvard University Press.

LLM Usage Statement

We used Claude (Anthropic) through Claude Code to write and test the experiment code, run and monitor the sessions, write the analysis scripts, and draft this report from our lab notes. The agents in the experiment are DeepSeek V4.1 Flash; the full-information judge is GPT-6 Luna. All numbers in this report are taken from the analysis outputs (results/sweep_urgency_pooled and results/orderbook_leak).

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