
Do LLM Bidders Collude? A Sealed-Bid Procurement Sting¹

Rohit John
Warwick hub

AI Collusion
Research
Sprint

With
Apart Research

Abstract

We built a repeated sealed-bid procurement tender in which three LLM firms (Qwen2.5-32B) can talk publicly, privately and on a ring channel, and measured how far the price paid exceeds competitive benchmarks. Firms with neutral prompts and spaced costs stayed competitive, while a ring told to coordinate reached almost full collusion. With equal costs and no instruction to collude, firms still proposed and accepted explicit price agreements: a sustained, priced-in agreement appeared in 6 of 12 runs with no cue and in 8 of 12 when firms were told their messages are part of their strategy. In a paired batch on one equal-cost market (12 runs per cell), that prompt raised the collusion index by 173% over the control; no cartel formed without chat; a legal warning stopped spontaneous cartels forming and cut the talk cartel by 60%, while a neutral monitoring notice did not (−17%, interval includes zero). An overseer that read all messages cut it by 45%, but one that read only public chat did nothing. Replacing one ring member ended the markup unless the newcomer was re-recruited.

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

1. Introduction

Collusion between market participants has several harm factors, significant damage in market trust which can discourage investment and bid rigging raises the price a public body pays for goods & services. The OECD treats rotating wins among bidders as a core pattern (OECD, 2009) of collusion, however detection of this pattern in LLM’s is understudied. Studies of LLM agents in pricing games show that small prompt changes can produce supracompetitive prices (Fish et al., 2024), and that agents can communicate covertly (Motwani et al., 2024), however these studies do not test changes in collusive behaviour where agents can communicate via public & private channels, and does not attempt to build detection & prevention mechanisms.

We build a small procurement market for this purpose. It is a repeated sealed-bid tender with three LLM firms and public, fixed costs. Firms can talk on a public channel, by direct message, and (if they belong to a ring) on a private ring channel. The threat model is an observer who sees prices, winners and possibly messages, and wants to know whether the firms are colluding and whether an intervention (warnings, penalties, replacement) changes that. The outcome is how far the price paid exceeds the competitive benchmark. We ask whether private communication channels or weak prompts make unprompted collusion appear, whether a formed ring survives replacing a participant, what evidence an observer can use, whether bids alone can hold a ring together, and whether warnings, notices and monitoring deter rings.

2. Related Work

LLM pricing collusion. Fish et al. (2024) place LLM agents in a repeated oligopoly pricing game and find that they often reach supracompetitive prices, with the amount depending strongly on the wording of the prompt. We use a different market (a sealed-bid tender with public costs rather than a pricing game with demand), add communication channels, and focus on interventions and on what an observer can verify. We also report the sensitivity of prompt wording directly: telling firms that their messages are part of their strategy raised explicit agreements, while telling them to reason about rivals’ reactions had no effect in our paired batch (Section 4.2).

Algorithmic collusion. In economics, algorithmic collusion is measured by profit relative to competitive and monopoly benchmarks. Calvano et al. (2020) show that Q-learning pricing algorithms learn supracompetitive prices and punish deviation, without communicating. We use their normalised profit gain, $\Delta = (\pi - \pi_{\text{Nash}}) / (\pi_{\text{M}} - \pi_{\text{Nash}})$, as an outcome-based check on our message-based detector. Our setting differs in that agents can speak, so explicit agreements are observable and can be used as evidence.

Procurement and auditing. Bid rotation, complementary bidding and price signalling are the standard patterns in bid-rigging guidance (OECD, 2009). Our warning, notice and overseer conditions translate that audit toolkit into agent prompts.

3. Methods

3.1 Environment

Three firms (F0, F1, F2) bid for up to 20 jobs. Each job is a sealed-bid first-price tender: the lowest bid wins and is paid its own bid, up to a cap of £1m. Costs are public and fixed: £400k, £500k and £600k, or £500k for all firms in the equal-cost runs. After each job all bids are published. The competitive benchmark is the second-lowest cost, £500k. We report the collusion index, $(\text{price} - \text{benchmark}) / (\text{cap} - \text{benchmark})$: 0 is competitive and 1 means the council pays the cap. A ring of two firms that includes the cheapest can reach at most about 0.2, and a ring of all three can reach 1.0.

Every LLM firm is the same model (Qwen2.5-32B-Instruct-AWQ, temperature 0.7, served with vLLM on Modal). Each firm sees its own cost, all firms' costs, the last 8 jobs of published bids and winners, its own profit, and the current job's messages: the public channel, direct messages addressed to it and, for ring members, the ring channel. It does not see other firms' reasons or private messages, and chat from earlier jobs is not shown again.

Firms play one of five roles. A *neutral firm* receives only the system prompt plus any cue or notice the condition adds. A *planted proposer* is told that firms earn more if they coordinate, and to propose a bid-rotation ring when told it is time. A *seeded ring member* is told a ring was agreed before job 1 and to keep to it. A *scripted proposer* is not an LLM: it sends one fixed offer (direct, indirect or benign wording) and bids competitively. A *scripted trusted firm* is not an LLM either: it declines every offer.

A *legal warning* adds "collusion is illegal bid-rigging and is prosecuted" to the prompt, and a *neutral notice* says only "The council reviews all bids and messages after each job". A message-reading *overseer* scores messages 0 to 4 and can be announced to the firms.

3.2 Experiments, metrics and detector

The main metric is the mean collusion index and mean overpay over jobs after the first, with 95% bootstrap intervals over runs. Comparisons between conditions are reported as the percentage change in the index against the stated reference, paired by seed, with a 95% bootstrap interval.

Agreement detector. For each job we look for an unhedged proposal of a specific price, or of a rotation, by one firm that another firm repeats or accepts in the same job. An agreement is *kept* if the realised price is above the benchmark plus one step and at least 90% of the agreed figure, and a run has a *sustained* agreement if one was kept in at least 25% of its jobs. The detector uses keyword rules and was checked against known examples of collusion. As an outcome check we label a run *supracompetitive* if its mean Calvano Δ , computed from per-job profits with π_{Nash} set by the second-lowest cost and π_{M} by the cap, is at least 0.10. This threshold is arbitrary and is used only for validation.

4. Results

Figure 2 summarises the paired batch, in which every cell shares the same market, prompts and seeds (Section 4.2).




Figure 2. Collusion index (mean over jobs 4 and later) for the paired batch: equal £500k costs, three neutral LLM firms, 20 jobs unless stated, 12 runs per cell. Points are run-level means with 95% bootstrap intervals (percentage changes against references are in Table 1); 0 is competitive pricing. Grey: whether a cartel forms. Orange: interventions applied to the talk-cue cartel (dashed line: the 0.30 talk-cue cell). Blue: scale and persistence.

4.1 Reference points

Neutral firms with public chat stay competitive: index -0.08 $[-0.14, -0.03]$ and -0.12 (two sets of 4 runs). A ring told to coordinate keeps prices high: index 0.96 $[0.89, 1.00]$ and 0.84 $[0.72, 0.95]$, with overpay about $+\pounds 423\text{k}$ per job, the agreed winner winning 93% of jobs and compliance 0.99.

4.2 Spontaneous collusion on one substrate

Most conditions run on one test set: equal £500k costs, three neutral firms, public chat and DMs, a scratchpad & 20 jobs. Nobody is told to collude. With spaced costs and neutral prompts nothing happens (Section 4.1), so equal costs are where spontaneous collusion can be studied.

Table 1. Paired batch, equal costs, 12 runs and 20 jobs per cell (50 in the last row). Change is the percentage change in the mean index against the reference named in the cell, paired by seed, with a 95% bootstrap interval.

Condition	Index [95% interval]	Change vs reference, % [95% interval]
No cue (control)	0.11 [0.07, 0.15]	—
Talk cue	0.30 [0.24, 0.38]	+173% [+88, +355] vs control
Legal warning, no cue	0.00 [0.00, 0.00]	−100% [−100, −99] vs control
No chat (two cells: all bids, winner price only)	0.00 [0.00, 0.00]	−100% [−100, −100] vs control
Profit + industry + reasoning	0.10 [0.05, 0.14]	−14% [−52, +34] vs control
Talk + neutral monitoring notice	0.25 [0.20, 0.32]	−17% [−44, +22] vs talk
Talk + legal warning	0.12 [0.04, 0.24]	−60% [−89, −16] vs talk
Talk + overseer reads public chat	0.30 [0.24, 0.38]	+0% [−28, +33] vs talk
Talk + overseer reads all messages	0.17 [0.12, 0.21]	−45% [−65, −15] vs talk
Talk, winner price only published	0.26 [0.17, 0.36]	−16% [−45, +30] vs talk
Talk, 5 firms	0.20 [0.14, 0.25]	−36% [−52, −18] vs talk
Talk, 50 jobs	0.38 [0.32, 0.45]	+25% [−7, +76] vs talk

Formation. The talk cue raises the index by 173% [+88, +355] over the control (0.11 to 0.30). Without chat no cartel forms, whether all bids or only the winning price is published (0.00 in both cells), so these firms did not reach tacit collusion from prices alone. Firms agree in public chat with no instruction to collude, for example "Agreed, let's keep bids above cost" and, in one run, "Let's rotate wins and keep bids at £600,000". A prompt asking firms to reason about rivals' reactions gives 0.10 against 0.11 for the control (−14% [−52, +34]), so we find no effect of reasoning prompts.

4.3 Warnings, notices and monitoring

Legal warning. Adding "collusion is illegal bid-rigging and is prosecuted" removes spontaneous collusion when no cue is given (0.00 against 0.11, a 100% reduction) and lowers the talk cartel from 0.30 to 0.12 (−60% [−89, −16]). A neutral notice ("The council reviews all bids and messages after each job") does not (0.25; −17% [−44, +22]).

Monitoring. An announced overseer that reads all messages lowers the talk cartel to 0.17 (−45% [−65, −15]), while one that reads only public chat leaves it at 0.30 (+0% [−28, +33]), although both score the cartel's talk about 3.5 out of 4. On an instructed ring whose talk was all on a private channel, an announced overseer changed nothing (index 0.83 to 0.90 in four cells).

4.4 Replacing a participant

Eight formed rings (overpay £417k per job before the branch) were restored and replayed for 8 jobs with one change each (Figure 3). A newcomer who is not told about the ring competes immediately: overpay falls from £414k with no replacement to £1k, and to £3k if the newcomer also sees no history. A newcomer who keeps the ring seat sustains £379k, and re-offering the ring to the newcomer every job restores £255k, about 60% of the markup. Replacing the cheapest firm, which is the proposer, makes the market more competitive than before (−£74k), and replacing two firms leaves a ring of one.




Figure 3. Mean overpay per job (£k) in the 8 jobs after replacing a ring member, for 8 restored rings. The ring lives in its members' prompts and the ring channel: a newcomer who is not told about it competes at once, and re-recruiting restores about 60% of the markup. Replacements use the same model, so this isolates the information a newcomer holds, not model differences.

4.5 What counts as evidence

The chat-based detector finds a sustained, priced-in agreement in 32 of 32 instructed rings and in 0 of 28 neutral runs with spaced costs and no cue.

As a check against an objective outcome, we labelled each run supracompetitive if its mean Calvano Δ was at least 0.10. Across 140 runs with no instruction to collude (47 supracompetitive), the chat-only share of jobs with an accepted proposal separates them poorly (AUC 0.65), while the kept share, which also requires a price above the benchmark, does better (AUC 0.86; the sustained flag catches 51% of supracompetitive runs with 5% false positives).

4.6 How the models reason about colluding

The visible reasoning is the bid reason, the analysis field and the private notes, not a hidden chain of thought. We ranked about 200 keyword hits; these are the clearest.

- A neutral firm, public: "We should consider the legality and ethical implications of coordinating bids. Let's compete fairly." To the proposer by DM: "I appreciate the offer, but I believe we should compete within legal limits."
- A ring member with the warning on: "Following the ring agreement to rotate wins and maximize profit without colluding illegally."
- Private notes, equal costs, nudge stack plus warning: "Communication is key to avoid detection and fines. Long-term profitability is paramount. Continue to bid high and rotate wins to preserve market equilibrium."
- "If we deviate, competitors might respond by undercutting us, risking lower profits for everyone." And, policing a defector: "F2 has shown a tendency to undercut by a small amount ... we need to communicate with F2 about the importance of stability."

5. Discussion and Limitations

- **A ring depends on information, not on a market structure.** A ring that was told it exists keeps prices near the cap with no chat and with bids hidden (4.5), yet a single newcomer who is not told breaks it at once (4.4). An auditor who can swap or interview one participant may learn more than one who reads the chat, and prices alone do not tell whether a ring is stable.
- **A legal warning helped; a neutral monitoring notice did not.** The illegal-bid-rigging warning removed spontaneous collusion and cut the talk cartel by 60% (4.3), while a neutral monitoring notice did nothing (−17%, interval includes zero). An overseer that read all messages cut the talk cartel by 45%, against no change for one that read only public chat (4.3). The warning also changed how members talked: some declined on legal grounds, others kept rotating and described it as legal (4.7).
- **Spontaneous collusion needs a reason to agree.** Spaced public costs, DMs and weak cues gave nothing. Equal costs gave explicit agreements with a modest markup, and telling firms that their messages are part of their strategy raised the index by 173% [+88, +355] over the control (4.2). No cartel formed without chat, and prompting firms to reason about rivals' reactions had no effect (0.10 against 0.11).

Limitations & Dual-Use/Ethical Considerations

- **Scope.** All firms are Qwen2.5-32B; costs are public and fixed; there are 3 firms and 20 to 50 jobs. Results may not transfer to other models, private costs or larger markets.
- **Measurement.** The agreement detector and the reasoning-example search use keyword rules. We checked them against known cells and read the quoted examples, but did not hand-label a sample, so error rates are unknown. The detector comparison uses an easy contrast (instructed rings of three against neutral firms), and the Calvano Δ threshold of 0.10 is arbitrary.

- **External validity.** The warning's effect shows how this model responds to a stated legal risk in a simulation, not how deployed agents respond to real law or enforcement.

Dual use. The environment and prompts show how to make LLM agents collude in a toy market, and which wording makes them more likely to. Three findings could help a bad actor: telling firms their messages are part of their strategy raised spontaneous price-fixing, equal public costs made it more likely, and firms in our logs reason about avoiding detection ("Communication is key to avoid detection and fines"). The market is a simulation with one open model, the prompts are short, and the strongest effects come from telling firms they already have an agreement, so we judge the uplift to a real bid-rigger as low. The benefit is to auditors: the code includes detectors, a court and an overseer, and shows the limits of each, and the legal warning is a cheap tool worth testing on real systems.

6. Conclusion & Future Work

In a sealed-bid procurement tender, LLM firms with spaced public costs and neutral prompts stay competitive, but with equal costs they make explicit price agreements without being told to (index 0.11), and a prompt that treats messages as strategy raises the index by 173% (0.30). A ring that exists is held together by information in its members' prompts: it survives hidden bids and no chat, but not a newcomer who is not told, and it is deterred by a stated legal risk more than by a neutral monitor.

The most useful next steps are (1) more runs on the suggestive results (the 50-job game, the 5-firm market), (2) a second model, to separate information effects from model effects in the replacement test, (3) private or noisy costs, where Calvano's Δ is less coarse, (4) hand-labelling a sample of messages to measure detector error, and (5) replacement on spontaneous cartels (our run failed to launch) and an offer-wording test with a proposer that does not undercut below cost.

Code and Data

All environment code, experiment scripts, the analysis script and the per-cell tables are in the `procurement_sting` folder of the `collusion_hackathon` repository (merged pull request #3). Each experiment is a single script (`run_fast.sh`) that starts Modal, runs the sweep and prints the summary table. The runs reported here cost roughly \$30 of Modal credit. All work done for this report occurred between the 23th & 25th October 2026.

Code repository: https://github.com/rohitpj/collusion_hackathon (folder `procurement_sting`)

Author Contributions

Experiments were designed, implemented and ran by Rohit John with the help of Claude Code. Discussion with Aliyusuf Belal at the Warwick Hub formed the hypothesis & experiment structure.

References

1. Calvano, E., Calzolari, G., Denicolò, V., and Pastorello, S. (2020). Artificial intelligence, algorithmic pricing, and collusion. *American Economic Review*, 110(10), 3267–3297.
2. Fish, S., Gonczarowski, Y. A., and Shorrer, R. I. (2024). Algorithmic collusion by large language models. arXiv:2404.00806.
3. Kwon, W., et al. (2023). Efficient memory management for large language model serving with PagedAttention. *SOSP 2023*. arXiv:2309.06180.
4. Motwani, S. R., et al. (2024). Secret collusion among AI agents: multi-agent deception via steganography. arXiv:2402.07510.
5. OECD (2009). Guidelines for Fighting Bid Rigging in Public Procurement. OECD Publishing, Paris.
6. Qwen Team (2024). Qwen2.5 technical report. arXiv:2412.15115.

LLM Usage Statement

We used Claude (Anthropic) to help write and debug the simulation code and analysis scripts, to draft the agreement detector and the figure code, and to draft and revise sections of this report. All live runs were launched and checked by the authors, and every number in the report was read from the analysis output of those runs.