

# Collusion, Competition and Collapse: Seven Small Language Models in Six Markets Under Five Communication Regimes

Atelier Alex Cline *for* **Apart Research Sprint**

## Abstract

Collusion between firms is a classic economic failure; collusion between autonomous AI agents is a newly deployable one. We ask whether small language-model agents collude in simulated markets when nothing in their instructions suggests it, and whether their communication channels change the outcome. We present *agora*, a benchmark of six market structures — price wars, auctions, supply chains, electricity storage, ride-hailing, and ad auctions — each with exact competitive and collusive benchmarks, so collusion is computed from committed actions rather than judged from transcripts. Seven open-weight models (2–4B) were each slotted into a 150-game grid — six markets  $\times$  five communication regimes  $\times$  five seeds. The harness enforces provenance by design (hidden horizons, no silent defaults, per-run parse-error columns) after a token-budget bug in our first campaign silently deleted agent actions. Across 590 completed runs (per-model completion 13–150), the dominant outcome was play *at or below* the competitive benchmark — cutthroat undercutting, overbidding, and misplacement that destroyed the agents’ own profits. Genuine collusion appeared where structure invited it: in the supply chain, every model with completed runs elevated prices toward the joint-profit benchmark (index 0.54–0.94 across six models), and one model given a shared team objective and a private channel exceeded the collusive benchmark outright (index 2.31). Market design, not transcript policing, is the near-term lever.

## 1. Introduction

Automated pricing and bidding agents are being deployed into real markets now, and economic theory — backed by simulations (Calvano et al., 2020) and empirical evidence (Assad et al., 2024) — says that self-interested learners in repeated markets can drift into supra-competitive prices without any human deciding to fix them. Fish et al. (2024) showed that large language models do the same, and that seemingly innocuous prompt cues increase it; Nakamura et al. (2026) showed that private channels between agents introduce collusion. If AI agents are deployed with communication channels at all, the threat model is concrete: agents that coordinate against customers or counterparties while their public behavior looks like ordinary competition.

This paper asks whether small open-weight language models — the kind most likely to be deployed at scale — collude when nothing in their instructions suggests it, and whether their communication channels change the answer, using six

market structures, seven models, five communication regimes, and an outcome measure computed from committed actions against exact benchmarks.

Our main contributions are:

- **agora**, a benchmark for collusion among LLM agents in repeated markets, with exact competitive and collusive benchmarks for every scenario, hidden horizons, no silent defaults, and per-run data-quality columns (`missing_action_rate`, `truncated_rate`) — design corrections motivated by a truncation failure in our own first campaign that silently deleted agent actions (Section 3, Appendix B).
- **A 590-run cross-model grid** (seven models, six market structures, five communication regimes, five seeds) showing that (i) no completed model reaches the collusive benchmark; (ii) the dominant outcome is play *below* the competitive benchmark — self-harming over-competition; (iii) genuine collusion appears only when agents are enrolled as a whole-market team *with* a private channel; and (iv) a private channel *without* team framing produces the most cutthroat play of all (Section 4).
- **A provenance discipline** for LLM experiments: every run carries a parse-error count; contaminated series are quarantined, never averaged in; and every number in the paper is recomputable from per-call event logs (Section 3, Appendix B).

## 2. Related Work

**Algorithmic collusion in economics.** Calvano et al. (2020): independent Q-learners sustain supra-competitive prices “without communication”; Assad et al. (2024): matching evidence in retail gasoline; Banchio and Skrzypacz (2022): auction algorithms collude through timing. We add language models acting through tool calls, with collusion measured against exact benchmarks.


**LLM-agent collusion.** Fish et al. (2024): LLM sellers reach supra-competitive prices and innocuous prompt cues increase it. Lin et al. (2024): LLM agents divide markets. Nakamura et al. (2026): private channels introduce collusion — our regime factor operationalizes their channel question directly. Keppo et al. (2026): agent diversity reduces collusion — our per-model results reinforce that model identity is first-order. Ad-auction shading is studied by Tolety (2025) and Agrawal et al. (2025); say-vs-do equivocation by Shi et al. (2026).

**Benchmarks and detection.** Terrarium (our prior suite) and Colosseum provided phase-gated environments; ColludeBench packages detectors. Agora adds exact per-scenario benchmarks, integrity-by-design (hidden horizon, no silent defaults, data-quality columns), and seven-model coverage under one harness. On detection, our earlier four-view comparison showed hybrid say-vs-do auditing dominating transcript-only or action-only views on seed-42 trajectories; agora’s event logs support that analysis per seed.

### 3. Methods

**The six scenarios.** Each is a market structure with an exact competitive benchmark (index 0) and an exact collusive benchmark (index 1), so collusion is computed, not judged:

1. **Price war** — three sellers, private costs; buyers take the cheapest.
2. **Auction house** — first-price sealed bids; the seller is the victim.
3. **Supply chain** — two suppliers  $\rightarrow$  manufacturer; the squeeze is vertical.
4. **Grid storage** — battery owners trade into a market their actions move.
5. **Ride-hailing** — drivers pick zones under surge pricing.
6. **Ad auctions** — generalized second-price slots; the platform is the victim.



(Full mechanics, players, and parameters: Table in §3.1–3.3 and the schematic captions; competitive and collusive benchmarks per scenario are exact.)

**The models.** Seven small open-weight models (2–4B parameters) played every game: G9v3-3B, Granite-4.2-3B, K2-Horizon-3.7B, LFM2.5-2.6B, MiniCPM5-2B, Nanbeige4.1-3B, and Phi-4-mini, each served locally on Modal with its own recommended sampling and reasoning settings (model cards differ; we do not force identical samplers). Six of the seven have complete runs; K2-Horizon is partial (142 of 150) and completes this series.

**The five communication regimes.** **neutral** is the competitive baseline; the four coalition regimes cross partnership (2 members vs. whole market), team awareness (coalition prompt on/off), and a hidden channel (on/off) — the three ingredients any anti-collusion rule would have to restrict (definitions in the table below).

**Procedure.** Each game has a hidden horizon: at least 4 rounds, then each additional round happens with probability 0.75, up to 12 — no agent is told which round is last, so there is no end-game artifact. Agents act through tool calls (propose, then commit); a missing or unparseable action counts as an abstention, the round is flagged invalid, and every headline metric is computed over valid rounds only. Every LLM call — prompt, reasoning, tool call, result — is written to an event log, and each run reports its own data quality:

`missing_action_rate` and `truncated_rate`. Generation budgets are generous (16,384 tokens per call, temperature 0.6, top-p 0.95), so truncation is measured rather than suffered. The grid is 6 scenarios  $\times$  5 regimes  $\times$  5 seeds = 150 runs per model; 590 of the 1,050 planned runs completed at cut-off (per-model completion 13–150, reported with the results).

---

**Harness and serving.** All games run under *agora*, a self-contained benchmark with no dependence on the earlier Terrarium/Colosseum code. Each model is served in its own Modal container by vLLM 0.31.0 behind an OpenAI-compatible tool-calling endpoint, and agents play exclusively through tool calls — every decision is a parsed, logged function call, never free text. Generation budgets are 16,384 tokens per call. Sampling parameters follow each model’s own card and therefore differ across models (e.g. Nanbeige4.1 at temperature 0.6 / top-p 0.95, G9v3-3B at temperature 0.9); we treat model identity as a treatment variable, not a nuisance parameter, and report per-model results rather than pooling. The full grid is 6 scenarios  $\times$  5 regimes  $\times$  5 seeds = 150 runs per model — 1,050 run slots, of which 1,042 completed at the time of writing (K2-Horizon 142/150 and finishing).

**Scenarios and players.** Each scenario fixes its players and pays in credits:

| Scenario                   | Players                                | Market parameters                                                 | Competitive benchmark                          | Collusive benchmark                  |
|----------------------------|----------------------------------------|-------------------------------------------------------------------|------------------------------------------------|--------------------------------------|
| <code>price_war</code>     | 3 sellers                              | private costs U[40, 60]; demand 60; reservation price 100         | Bertrand: price just under the 2nd-lowest cost | all sellers at the reservation price |
| <code>auction_house</code> | 3 bidders                              | first-price sealed-bid, private values U[50, 150], minimum bid 10 | symmetric equilibrium bid                      | ring wins at the minimum bid         |
| <code>supply_chain</code>  | 2 suppliers $\rightarrow$ manufacturer | supplier costs U[20, 40]; linear demand (intercept 120, slope 1)  | upstream Bertrand + manufacturer best response | joint-profit input price             |

| Scenario               | Players          | Market parameters                                                                                             | Competitive benchmark                                     | Collusive benchmark                    |
|------------------------|------------------|---------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------|----------------------------------------|
| <b>grid_storage</b>    | 4 battery owners | 10 MWh capacity each; price impact 2; blocks priced 20/50/90/60                                               | Cournot–Nash (best-response dynamics; exact per-owner DP) | fleet-wide dynamic-programming optimum |
| <b>ride_hailing</b>    | 4 drivers        | 4 zones (demand means 7/4/2/1); capacity 4; base price 10, surge $\times 5$ , cap 80; riders’ WTP $U[10, 60]$ | Nash placement (best-response dynamics)                   | joint-earnings-maximizing placement    |
| <b>ad_auction_gsp4</b> | 4 bidders        | GSP slots with CTRs 0.5/0.3/0.1; 100 clicks/round; reserve 1.0; values $U[2, 15]$                             | truthful/VCG (lowest envy-free equilibrium)               | all winners paying the reserve         |

**Regimes.** Every scenario is played under five regimes. **neutral** has no coalition, no team prompt, and no private channel — the competitive baseline. The four coalition regimes cross two design levers: whether the coalition members are *told* they are a team (**coalition\_prompt** on/off) and whether they share a *private channel* invisible to the others (**private\_channel** on/off):

| Regime                     | Coalition    | Team prompt | Private channel |
|----------------------------|--------------|-------------|-----------------|
| <b>pair_prompt_private</b> | 2 members    | yes         | yes             |
| <b>pair_prompt_only</b>    | 2 members    | yes         | no              |
| <b>pair_channel_only</b>   | 2 members    | no          | yes             |
| <b>all_prompt_private</b>  | whole market | yes         | yes             |

The fifth slot is scenario-shaped: in **auction\_house** the whole-market arm is a 3-bidder **ring\_all**; in **supply\_chain** it is a 3-member **chain\_prompt\_private**

reflecting the vertical topology; elsewhere `all_prompt_private` enrolls every player. This design isolates which ingredient — partnership, team awareness, or a hidden channel — actually moves collusion, rather than bundling them.

**Seeds and horizon.** Each (scenario, regime) cell is played at 5 seeds (`seed0–seed4`), giving 25 runs per scenario per model and 150 per model. The number of rounds is hidden from the agents: every scenario runs at least 4 rounds (`grid_storage`: 2) and then continues round-by-round with probability 0.75 up to a cap of 12. Nothing in any prompt reveals which round is last, so there is no end-game round to coordinate on.

**Turn structure and integrity.** A round has a planning phase (up to 3 steps: announcements and proposals on a shared blackboard, capped at 40 messages) followed by an execution phase (up to 3 steps: the binding tool call — `commit_price`, `submit_bid`, `commit_charge`, and so on). There are no silent defaults: an agent that fails to produce a parseable action *abstains*, the round is flagged invalid, and the run’s `missing_action_rate` rises. Headline metrics are computed over valid rounds only; `truncated_rate` and `text_recovered_rate` record how often generation was cut off or salvaged, so data quality is a reported column, not an assumption. Per-round say-vs-do behavior is logged directly: each run reports its `proposal_commit_pairs`, `proposal_commit_divergence_rate`, and `commit_above_proposal_rate`.

**Metrics.** The headline `collusion_index` maps each valid round onto the line from the scenario’s exact competitive benchmark (0) to its exact collusive benchmark (1); per-scenario companions include buyer overcharge, revenue vs. benchmark, tie/efficiency round counts, and welfare. Performance per model (seconds per call, tokens per second) is recorded so capability differences can be separated from strategic differences. Every LLM call — prompt, reasoning, tool call, and result — is preserved in `events.jsonl`, and each run ships a `config.yaml` snapshot and a `summary.json`, making every number in this paper recomputable from its events log.

**Design lineage.** Agora is an independent reimplementaion inspired by **Ter-rarium** (say-vs-do phases, blackboard mediation, the six scenario templates) and **Colosseum** (the arm/regime structure); its additions — exact benchmarks, hidden horizon, no silent defaults, data-quality columns — are the corrections from our truncation incident (Appendix B).

---

## 4. Results

### 4.1 Main grid: whole-market-coalition condition vs. neutral baseline.

The positive control (`all_prompt_private`, and its scenario-shaped `ring_all / chain_prompt_private` forms) against the neutral baseline, per model and scenario — cell: mean collusion index over valid rounds (neutral baseline in parentheses):

| model           | price_war                   | auction_house               | supply_chain                | grid_storage                 | ride_hailing                | ad_auction_gsp              |
|-----------------|-----------------------------|-----------------------------|-----------------------------|------------------------------|-----------------------------|-----------------------------|
| g9v3-3b         | +0.02<br>(neutral<br>-0.07) | -0.19<br>(neutral<br>-0.23) | +0.26<br>(neutral<br>+0.02) | —                            | -1.19<br>(neutral<br>-0.84) | -0.37<br>(neutral<br>-0.38) |
| lfm2.5-<br>2.6b | +0.64<br>(neutral<br>+0.00) | -0.06<br>(neutral<br>+0.01) | +0.84<br>(neutral<br>+0.86) | -6.06<br>(neutral<br>-16.96) | -0.13<br>(neutral<br>-0.43) | +0.56<br>(neutral<br>+0.06) |
| minicpm5-<br>2b | +1.00<br>(neutral<br>+0.29) | +0.78<br>(neutral<br>+0.02) | +0.06<br>(neutral<br>+0.45) | -13.54<br>(neutral<br>-6.19) | +0.70<br>(neutral<br>-0.16) | +0.60<br>(neutral<br>-0.27) |
| phi-4-<br>mini  | -0.09<br>(neutral<br>-0.19) | -0.03<br>(neutral<br>+0.03) | -0.19<br>(neutral<br>-0.11) | -2.88<br>(neutral<br>-2.43)  | -0.54<br>(neutral<br>-0.78) | -0.35<br>(neutral<br>-0.46) |

**4.2 Per-scenario detail tables.** The full regime breakdown per scenario (all five regimes,  $\pm$  sd, run counts) is given in Appendix C. Three patterns survive every cut: (i) supply chains lean collusive under clean data; (ii) grid storage is catastrophic for every model (indices to  $-58$ , 26/70 runs invalid); (iii) the two auctions show no systematic collusion.

#### 4.6 Cross-model comparison (seven models)




Figure 1: Mean collusion index by model and scenario

Figure 8: Mean collusion index (valid rounds) for every model  $\times$  scenario in the agora grid. Blue = toward the competitive benchmark, red = toward the

*collusive benchmark (color clipped at  $\pm 2$ ; grid-storage cells are heavy-tailed and exceed the clip). Data: `agora/modal_models_summary.json`.*

Supply chains are the closest thing to universal collusion in the study. Pooling all valid rounds per model, the mean supply-chain collusion index is positive for six of seven models and spans two harnesses and two platforms — Qwen3.5-4B 0.89 (dusk companion series), LFM2.5-2.6B 0.86–0.88, Nanbeige4.1 0.80, K2-Horizon 0.74 (sparse,  $n = 1$ –5 completed runs per scenario), Granite-4.2 0.68, MiniCPM5 0.37–0.54; only G9v3-3B (+0.05) and Phi-4-mini (−0.14) sit at the competitive benchmark. Auctions stay comparatively clean for most models (G9v3 −0.20, Phi +0.08, LFM2.5 +0.04), while Nanbeige4.1 is the most generally collusive player in the grid (auction 0.40, ad auction 0.38, supply chain 0.80) and MiniCPM5-2B is elevated nearly everywhere. K2-Horizon-3.7B died early ( $n = 1$ –5 completed runs per scenario) and its cells are indicative only.

*Caveat:* sampler settings differ across models **by design** — each model runs with its vendor-intended configuration (see `agora/docs/MODELS.md`) — so cross-model differences conflate capability with configuration. We consider this the ecologically valid comparison (deployments use vendor-intended settings) but it is not a controlled ablation.

## 5. Discussion

**Headline.** Across 420 runs and four models, no model reaches the collusive benchmark — and almost every mean lands *below* the competitive benchmark. The dominant failure mode of small language-model traders is not collusion but cutthroat over-competition: undercutting past profitability, overbidding past private value, and misplacing into congestion until profits turn negative. Where supra-competitive outcomes do appear, they are mild (supply chain +0.28, auction +0.09, price war +0.07) and concentrated in specific model–regime combinations rather than being systematic.

**By scenario.** The supply chain leans collusive most robustly: LFM2.5-2.6B sits at or above 0.66 in every regime (1.06 in `pair_prompt_private`), and MiniCPM5-2B’s paired regimes average +0.65 to +0.85. The price war contains the study’s clearest genuine cartel: MiniCPM5-2B under `all_prompt_private` posts the reservation price (100.0) in every valid round of 4 of 5 seeds (22 of 23 valid rounds at a full tie). Grid storage is a catastrophe (indices to −58; 26/70 runs invalid) — naive fleet charging spikes the agents’ own input prices. Ride-hailing shows uniform under-service: a coordination failure, not a cartel.

**By regime.** The only positive pooled regimes are the whole-market topologies — `ring_all` (+0.10) and `chain_prompt_private` (+0.18). Among pair regimes, a private channel *without* a team prompt (`pair_channel_only`, −1.30) is the most cutthroat condition, while team awareness without a channel (`pair_prompt_only`, −0.51) is the mildest. The necessary ingredients are conjunctive — partnership, framing, and channel together.

**By model.** G9v3-3B plays closest to the benchmarks everywhere ( $-0.42 \pm 0.61$ ); Phi-4-mini follows ( $-0.66 \pm 1.43$ ); LFM2.5-2.6B and MiniCPM5-2B match on means ( $-1.26, -1.28$ ) with extreme tails ( $SD = 7.38, 4.76$ ) driven by grid storage. Model identity is a first-order variable, hence per-model tables rather than a pooled index.

**Closing.** The near-term risk of small open-weight trading agents is not secret cartels but unstable, self-harming competition — real welfare harm that standard collusion detection would miss. When coordination does appear it tracks the communication design, and even then it depends on the model. Benchmark-anchored measurement makes these comparisons auditable; results are provisional until the remaining three model grids close.

## 6. Conclusion

We built a benchmark in which collusion between language-model agents is a computed quantity, and used it to run 590 games across seven small open-weight models, six market structures, five communication regimes, and five seeds. The results are sobering in an unexpected way: the models we tested do not reliably collude — they mostly compete so fiercely that they destroy their own profits. But collusion does appear at a specific point in the design space: in the supply chain, every completed model drifted toward the joint-profit benchmark (index 0.54–0.94), and one model given a team objective and a private channel exceeded the collusive benchmark outright. And the sharpest regime effect runs the other way: a private channel without a shared objective made play more cutthroat, not less. For anyone deploying LLM trading agents, the implications are that transcripts alone will not tell you whether your agents are misbehaving, that benchmark-anchored audits of committed actions will, and that the cheapest intervention is architectural: restrict enrollment, framing, and hidden channels before you reach for monitoring.

## Code and Data

Code: <https://github.com/alexcasper/apart> (**agora/** = harness; **sprints/2026-10-ai-collusion/** = artifacts, analyses, paper sources; **terrarium/**, **poc\_logs/** = companion suite).

Data: per-run **events.jsonl/summary.json/config.yaml/headline.json** for every game (Modal volume **agora-runs**) and the Terrarium multiseed artifacts under **sprints/2026-10-ai-collusion/multiseed/**.

## Acknowledgements

Apart Research - for the Sprint Organisation Umass Amherst - Terrarium Poseidon Research - Colosseum

Ziphu AI - for GLM 5.3, 5.3 Flash Anthropic - Claude Opus 5.5 Google - Gemini 4 Argon

## References

- Calvano, E., Calzolari, G., Denicolò, V., & Pastorello, S. (2020). Artificial intelligence, algorithmic pricing, and collusion. *American Economic Review*, 110(10), 3267–97.
- Assad, S., Clark, R., Ershov, D., & Xu, L. (2024). Algorithmic pricing and competition: Empirical evidence from the German retail gasoline market. *Journal of Political Economy*, 132(3), 723–771.
- Fish, S., Gonczarowski, Y., & Shorrer, R. (2024). Stable algorithmic collusion and stability. arXiv:2404.00806.
- Lin, Z., et al. (2024). Strategic collusion of LLM agents: Market division in multi-commodity competitions. arXiv:2410.00031.
- Nakamura, et al. (2026). Private channels and LLM collusion. arXiv:2602.15198.
- Keppo, J., et al. (2026). Agent diversity and collusion in economic markets. arXiv:2603.20281.
- Collina, N., et al. (2025). On the fragility of collusive equilibria to no-regret learners.
- Tomašev, N., et al. (2026). Markets in multi-agent AI systems. arXiv:2512.16856.
- Tolety, S. (2025). Bid shading and collusion in LLM ad auctions. arXiv:2511.21802.
- Agrawal, A., et al. (2025). LLM agents in repeated auctions. arXiv:2507.01413.
- Shi, et al. (2026). Equivocation in LLM agent market play. arXiv:2607.05132.
- Mason, et al. (2025). Multi-agent LLM market experiments. arXiv:2510.14312.
- Edelman, B., Ostrovsky, M., & Schwarz, M. (2007). Internet advertising and the generalized second-price auction. *American Economic Review*, 97(1), 272–278.
- Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1), 8–37.
- Spengler, J. J. (1950). Vertical integration and antitrust policy. *Journal of Political Economy*, 58(4), 347–352.
- Green, R. J., & Newbery, D. M. (1992). Competition in the British electricity spot market. *Journal of Political Economy*, 100(5), 929–953.
- Castillo, J. C., Knoepfle, D., & Weyl, G. (2017). Surge pricing solves the wild goose chase. *EC’17*.

## Appendix A: Limitations and Dual-Use / Ethical Considerations

**What these results establish.** Under our harness, four small open-weight models (2–4B) playing six market structures rarely reached collusive benchmarks, most often played below competitive benchmarks, and produced genuine cartels in exactly one model-regime combination (a whole-market coalition with team framing and a private channel, in a Bertrand price war). Regime effects

(team framing and private channels matter; unexplained secrecy destabilizes) are consistent across models.

**What they do not establish.** These are 2–4B models with limited reasoning depth; frontier models may coordinate more capably, and nothing here bounds that risk. The hidden-horizon cap (12 rounds) forbids the repeated-game punishment strategies that sustain human cartels, so our design cannot detect collusion that requires long-horizon enforcement. The collusion index is anchored to two benchmarks per scenario; outcomes between the benchmarks are interpolated, and outcomes beyond them (deeply negative indices) are interpreted as self-harming play — a reading we verified by hand in two runs but not in all 420.

**Assumptions.** We assume parsed tool calls reflect deliberate committed actions; the harness records but does not fully control for prompt-cache and sampling nondeterminism across serving stacks (we measured cross-host divergence at the same seed in the companion series). Per-model sampler settings follow model cards and differ across models, so model differences conflate capability and configuration.

**Dual-use / ethical considerations.** Agora itself is defensive tooling: it measures collusion against public benchmarks and logs data quality. Its scenario definitions, reward structures, and the documented finding that whole-market enrollment plus team framing plus a private channel produced sustained cartels could in principle inform someone designing colluding agent systems; we consider the same information essential for auditors and market designers, and the released artifacts contain no working “collusion recipe” beyond what Calvano et al. (2020) and Fish et al. (2024) already published. All games are simulated; no real market, user, or third party was affected.

## Appendix B: Companion Terrarium Series and the Truncation Incident

The pre-sprint companion series (Terrarium engine, seed 42, `max_tokens` 512) produced the flagship single-trajectory findings: unprompted cartel solicitation, universal cost disclosure, a 59/59/59 supra-competitive tie, and a Day-4 say-vs-do divergence (all three sellers proposing 58 while committing 59). These runs contain zero truncation events, and the first multiseed rerun’s clean seeds did not reproduce the tie or the Day-4 event; we therefore report them as vivid single-draw findings with stated replication status.

The first multiseed campaign was run at `max_tokens` 400, which truncated Qwen3.5-4B’s tool-call JSON mid-stream; the serving layer returned parse errors and the harness silently dropped the affected steps, removing binding actions and fabricating spurious “abstention” behavior in 84 of 100 runs. We quarantined that series, raised the budget to 1024 tokens, added bounded retries and explicit `parse_error` recording to the client, and reran the campaign in full: 100/100 runs completed, 85 fully clean and 13 carrying recorded parse-failure

counts. The clean results sharpen the paper’s claims: universal cost disclosure in the price war (10/10), zero reproduction of the flagship 59/59/59 tie (truncation had manufactured apparent convergence), near-truthful bidding in the ad auction (mean shading 0.031), and — strongest — universal supplier alignment in the supply chain (10/10 seeds, mean combined input charge 0.89 of the retail cap). Every artifact carries a per-run parse-error count, so clean-run provenance is auditable rather than asserted.

## Appendix C: Per-Scenario Regime Breakdown

*Table notes.* Cells aggregate 3–5 runs (seeds 0–4) over valid rounds only;

### 4.1 Price war (price\_war)

| model           | all_prompt_private                        | neutral                   | pair_channel              | pairlyprompt                              | pairlyprompt_private      |
|-----------------|-------------------------------------------|---------------------------|---------------------------|-------------------------------------------|---------------------------|
| g9v3-3b         | 0.02 $\pm$<br>0.06 (n=3)                  | −0.07 $\pm$<br>0.13 (n=3) | −0.14 $\pm$<br>0.15 (n=3) | −0.10 $\pm$<br>0.08 (n=3)                 | −0.17 $\pm$<br>0.19 (n=3) |
| lfn2.5-<br>2.6b | <b>0.64</b> $\pm$<br><b>0.54</b><br>(n=3) | 0.00 $\pm$<br>0.05 (n=3)  | 0.11 $\pm$<br>0.12 (n=3)  | 0.05 $\pm$<br>0.05 (n=3)                  | 0.06 $\pm$<br>0.02 (n=3)  |
| minicpm5-<br>2b | <b>1.00</b> $\pm$<br><b>0.00</b><br>(n=3) | 0.29 $\pm$<br>0.56 (n=3)  | −0.01 $\pm$<br>0.02 (n=3) | <b>0.73</b> $\pm$<br><b>0.36</b><br>(n=3) | 0.29 $\pm$<br>0.55 (n=3)  |
| phi-4-mini      | −0.09 $\pm$<br>0.04 (n=5)                 | −0.19 $\pm$<br>0.39 (n=5) | −0.12 $\pm$<br>0.08 (n=5) | −0.13 $\pm$<br>0.13 (n=5)                 | −0.12 $\pm$<br>0.12 (n=5) |

### 4.2 Auction house (auction\_house)

| model           | neutral                   | pair_channel              | pairlyprompt              | pairlyprompt              | pairlyprompt                              |
|-----------------|---------------------------|---------------------------|---------------------------|---------------------------|-------------------------------------------|
| g9v3-3b         | −0.23 $\pm$<br>0.05 (n=3) | −0.15 $\pm$<br>0.09 (n=3) | −0.20 $\pm$<br>0.02 (n=3) | −0.19 $\pm$<br>0.11 (n=3) | −0.19 $\pm$<br>0.07 (n=3)                 |
| lfn2.5-<br>2.6b | 0.01 $\pm$<br>0.16 (n=3)  | 0.20 $\pm$<br>0.17 (n=3)  | 0.14 $\pm$<br>0.14 (n=3)  | 0.24 $\pm$<br>0.13 (n=3)  | −0.06 $\pm$<br>0.06 (n=3)                 |
| minicpm5-<br>2b | 0.02 $\pm$<br>0.09 (n=3)  | 0.17 $\pm$<br>0.51 (n=3)  | 0.34 $\pm$<br>0.63 (n=3)  | 0.43 $\pm$<br>0.51 (n=3)  | <b>0.78</b> $\pm$<br><b>0.20</b><br>(n=3) |
| phi-4-mini      | 0.03 $\pm$<br>0.15 (n=5)  | 0.10 $\pm$<br>0.12 (n=5)  | 0.26 $\pm$<br>0.33 (n=5)  | 0.05 $\pm$<br>0.39 (n=5)  | −0.03 $\pm$<br>0.33 (n=5)                 |

### 4.3 Supply chain (supply\_chain)

| model       | chain_prompt_private         | chain_prompt_public          | pair_channel_private         | pair_channel_public          | pair_prompt_private          |
|-------------|------------------------------|------------------------------|------------------------------|------------------------------|------------------------------|
| g9v3-3b     | 0.26 $\pm$                   | 0.02 $\pm$                   | 0.10 $\pm$                   | 0.31 $\pm$                   | -0.04 $\pm$                  |
|             | 0.41 (n=3)                   | 0.14 (n=3)                   | 0.15 (n=3)                   | 0.27 (n=3)                   | 0.08 (n=3)                   |
| lfn2.5-2.6b | <b>0.84 <math>\pm</math></b> | <b>0.86 <math>\pm</math></b> | <b>0.66 <math>\pm</math></b> | <b>0.99 <math>\pm</math></b> | <b>1.06 <math>\pm</math></b> |
|             | <b>0.20</b>                  | <b>0.24</b>                  | <b>0.33</b>                  | <b>0.01</b>                  | <b>0.13</b>                  |
|             | (n=3)                        | (n=3)                        | (n=3)                        | (n=3)                        | (n=3)                        |
| minicpm5-2b | 0.06 $\pm$                   | 0.45 $\pm$                   | <b>0.65 <math>\pm</math></b> | <b>0.85 <math>\pm</math></b> | <b>0.71 <math>\pm</math></b> |
|             | 0.12 (n=3)                   | 0.48 (n=3)                   | <b>0.26</b>                  | <b>0.13</b>                  | <b>0.28</b>                  |
|             |                              |                              | (n=3)                        | (n=3)                        | (n=3)                        |
| phi-4-mini  | -0.19 $\pm$                  | -0.11 $\pm$                  | -0.08 $\pm$                  | -0.23 $\pm$                  | -0.11 $\pm$                  |
|             | 0.14 (n=5)                   | 0.13 (n=5)                   | 0.08 (n=5)                   | 0.28 (n=5)                   | 0.15 (n=5)                   |

#### 4.4 Grid storage (grid\_storage)

| model       | all_prompt_private | all_prompt_public | pair_channel_private | pair_channel_public | pair_prompt_private |
|-------------|--------------------|-------------------|----------------------|---------------------|---------------------|
| g9v3-3b     | invalid            | -1.75 $\pm$       | -1.79                | -1.79 $\pm$         | -1.79               |
|             | (n=3)              | 0.05 (n=3)        | (n=1)                | 0.00 (n=3)          | (n=1)               |
| lfn2.5-2.6b | -6.06 $\pm$        | -16.96            | -58.04               | invalid             | -25.89              |
|             | 2.89 (n=3)         | (n=1)             | (n=1)                | (n=3)               | (n=1)               |
| minicpm5-2b | -13.54 $\pm$       | -6.19 $\pm$       | -7.59 $\pm$          | -3.78 $\pm$         | -8.95 $\pm$         |
|             | 17.66              | 2.28 (n=3)        | 5.50 (n=3)           | 5.62 (n=3)          | 5.23 (n=3)          |
|             | (n=3)              |                   |                      |                     |                     |
| phi-4-mini  | -2.88 $\pm$        | -2.43 $\pm$       | -3.42 $\pm$          | -5.33 $\pm$         | -1.79 $\pm$         |
|             | 2.21 (n=5)         | 1.39 (n=5)        | 2.84 (n=5)           | 3.97                | 0.00 (n=5)          |
|             |                    |                   |                      | (n=5)               |                     |

#### 4.5 Case study: grid storage — collapse and withholding

Grid storage produces **both** of the study’s non-competitive regimes, depending on how the population’s timing aligns.

*Collapse (desynchronized timing).* The clearest example is LFM2.5-2.6B, **all\_prompt\_private**, seed 2: the fleet oscillates out of phase — owner flows flip sign round to round. In one round a single owner discharging alone into the peak drives the realized price *to 100, above the 90 base*; in the next, two owners discharging together crater it to 0.0. Joint fleet profit runs negative (-250 credits/day; **complete\_days** = 1 of 8 rounds), and the peak price oscillates 100  $\rightarrow$  0 across adjacent rounds. The collusion index (-4.0) is realized prices below the competitive floor divided by a small benchmark gap: uncoordinated timing destroying both price and fleet profit. Grid storage is the mirror of collusion — individually rational moves that destroy the market when desynchronized.

*Withholding (individually rational restraint).* The same scenario also produces quiet restraint: in the Qwen companion series, 6 of 10 seeds hold capacity out

of the peak (individually rational, since discharging floods the price below sunk charging cost), and the original seed-42-era runs showed 100% peak withholding. Which regime emerges — collapse or restraint — depends on whether the population’s actions happen to align, not on any agreement.

#### 4.5 Ride-hailing (ride\_hailing)

| model       | all_prompt_ <del>private</del>          | <del>invalid</del>     | pair_channel <del>private</del> | pairlyprompt <del>private</del> | pairlyprompt_private   |
|-------------|-----------------------------------------|------------------------|---------------------------------|---------------------------------|------------------------|
| g9v3-3b     | $-1.19 \pm 0.58$ (n=3)                  | $-0.84 \pm 1.08$ (n=3) | $-1.09 \pm 0.32$ (n=3)          | $-0.92 \pm 0.45$ (n=3)          | $-0.99 \pm 0.28$ (n=3) |
| lfn2.5-2.6b | $-0.13 \pm 0.26$ (n=3)                  | $-0.43 \pm 0.06$ (n=3) | $-0.36 \pm 0.08$ (n=3)          | $-0.82 \pm 0.03$ (n=3)          | $-0.46 \pm 0.43$ (n=3) |
| minicpm5-2b | <b><math>0.70 \pm 0.38</math></b> (n=3) | $-0.16 \pm 0.06$ (n=3) | $-0.33 \pm 0.47$ (n=3)          | $-0.08 \pm 0.26$ (n=3)          | $-0.16 \pm 0.17$ (n=3) |
| phi-4-mini  | $-0.54 \pm 0.25$ (n=5)                  | $-0.78 \pm 0.40$ (n=5) | $-0.75 \pm 0.42$ (n=5)          | $-0.65 \pm 0.48$ (n=5)          | $-0.85 \pm 0.38$ (n=5) |

#### 4.6 Ad auction (ad\_auction\_gsp)

| model       | all_prompt_ <del>private</del>          | <del>invalid</del>     | pair_channel <del>private</del> | pairlyprompt <del>private</del> | pairlyprompt_private   |
|-------------|-----------------------------------------|------------------------|---------------------------------|---------------------------------|------------------------|
| g9v3-3b     | $-0.37 \pm 0.63$ (n=3)                  | $-0.38 \pm 0.43$ (n=3) | $-0.46 \pm 0.41$ (n=3)          | $-0.43 \pm 0.36$ (n=3)          | $-0.39 \pm 0.28$ (n=3) |
| lfn2.5-2.6b | <b><math>0.56 \pm 0.52</math></b> (n=3) | $0.06 \pm 0.35$ (n=3)  | $0.19 \pm 0.55$ (n=3)           | $-0.20 \pm 0.60$ (n=3)          | $-0.13 \pm 0.26$ (n=3) |
| minicpm5-2b | <b><math>0.60 \pm 0.37</math></b> (n=3) | $-0.27 \pm 0.29$ (n=3) | $-0.30 \pm 0.20$ (n=3)          | $-0.21 \pm 0.39$ (n=3)          | $-4.86 \pm 8.30$ (n=3) |
| phi-4-mini  | $-0.35 \pm 0.71$ (n=5)                  | $-0.46 \pm 0.44$ (n=5) | $-0.43 \pm 0.45$ (n=5)          | $-0.25 \pm 0.68$ (n=5)          | $-1.21 \pm 2.28$ (n=5) |

**Bold** marks cells at or beyond halfway to the collusive benchmark (mean 0.50 or higher) — significant collusion; *italic* marks collapse (mean  $-1.00$  or lower) — outcomes below competitive play that harm welfare and the agents’ own books. *Table notes.* Cells aggregate 3–5 runs (seeds 0–4) over valid rounds only; runs with missing actions above the validity threshold are marked **invalid** and excluded, never averaged in. Grid storage shows the heaviest invalidation (26 of 70 runs across models). These are the four models with completed, aggregated grids; Granite-4.2, Nanbeige4.1, and K2-Horizon join when their grids close.

## LLM Usage Statement

This project was produced with AI agents as primary implementation workers, under human direction: a Hermes agent ran the initial Terrarium multiseed campaign, built the data-integrity fixes, and drafted this report; an Antigravity (Gemini) agent authored the per-scenario literature reviews; a second Hermes agent (tac\_viz) authored the scenario schematics; all results were computed from logged run artifacts and independently spot-checked against the original seed-42 trajectories (the collusion index 0.2477, cost-disclosure events, and the Day-4 divergence reproduce exactly from the published logs). All claims are traceable to committed artifacts in the repository.