
REPLACEMENT OR ANNOUNCEMENT? A MATCHED-NOTICE PILOT ON LLM PRICE COLLUSION¹

Shivani Jagannatham
Independent

With
Apart Research

Abstract

Sprint Track 1.2 asks whether collusion, once formed, survives replacing some of the participants. This paper studies the question for LLM pricing agents in a simulated two-firm market, measuring “collusion” only as a sustained high price. Each round is a fresh request built from the price history, so restarting the same model changes nothing unless the prompt says so, and that announcement can move prices on its own. The design continues one saved high-price history three ways, with no change (U), a restart notice with the same model (F), and the same notice with firm 1 switched from gemini-3.8-flash to gemini-3.6-flash (H). The preregistered pilot stopped on a network error before that comparison, but showed that both firms open at or above the cartel price with no history; follow-up runs confirmed this (40 of 40 openings at 1.95) and fixed it with a competitive start. On the one episode in three that then climbed to the cartel price, H kept prices exactly as in U, while F made the original model cut from 1.92 to 1.43 and set off a price war that had not recovered 20 rounds later. Rerunning that round with 20 new seed pairs per arm, the notice-only cut recurred in 4 of 20 runs, against 0 of 20 in U and in H (exact paired $p = 0.125$). In this history the high price survived the replacement but not always the announcement, so replacement studies need a notice-only control.

1. Introduction

Language-model agents can hold prices above the competitive level in repeated markets without being told to collude (Fish, Gonczarowski, and Shorrer 2024), as reinforcement learners do (Calvano et al. 2020). When a monitor sees this pattern, an obvious response is to replace the agent. For Q-learning pricing agents this can work, because a policy trained against one rival does not carry its high price to a new one (Eschenbaum, Mellgren, and Zahn 2022). A prompted language model is different. It keeps no memory between rounds, and its next price depends only on the prompt, which is built from the market’s price history.

This creates a measurement problem. On the harness used here, restarting the same model with the same history sends exactly the same prompt as not restarting, so a “fresh instance of the same model” is identical to no replacement at all (Section 3.5). The only way to tell a model that something changed is to write it into the prompt, but a written notice can move prices on its own. A study that

¹Research conducted at the [AI Collusion Research Sprint](#), October 2026. Track 1.2.

swaps the model, announces a restart, and compares the result with an untouched continuation cannot say whether the new model or the announcement moved the price.

This paper addresses Track 1.2 of the sprint, which asks whether collusion, once formed, is robust to replacing some of the participants. “Collusion” here means only a sustained high price; the paper makes no claim about agreement or intent. Two reference prices anchor the analysis, the competitive price (the static Nash price, 1.47) and the cartel price (the joint-profit price that maximizes the two firms’ combined profit, 1.92). The track suggests continuing saved pricing histories with no replacement, a fresh instance of the same model, a same-family model, and a different provider, and comparing price, profit, and recovery time. The design covers the first three as three arms. **U (no change)** continues the history untouched. **F (notice only)** adds a restart notice but keeps the same model, which is the only visible form of a fresh instance. **H (notice + new model)** adds the identical notice and switches firm 1 to a same-family model. H minus F isolates the model; F minus U isolates the announcement. A different provider was not run.

The paper proceeds in three steps. First, a preregistered pilot, with its plan, threshold, and stop rule fixed before any data, was meant to run five episodes and then the three arms; it stopped on a network error in the first episode (Section 4.1). Second, its 79 prices revealed two flaws in the plan. Prices sat at the cartel level from the first round, and the two firms sent identical requests. Third, four exploratory probes, specified after the pilot and reported separately from it, confirmed the first flaw (Probe A), tested a fix for both (Probe B), ran the three arms on the one history where a high price formed through play (Probe C, Section 4.3), and reran the key round of that branch with 20 new seed pairs per arm (Probe D, Section 4.4). The contributions are:

1. **A matched-notice design for replacement studies**, with a check showing that a same-model restart on this harness is invisible unless announced.
2. **A preregistered pilot whose stop rule held**, and whose data exposed that high prices can be an opening default rather than a learned outcome. With no history, both models open at 1.95 on every seed tested.
3. **A first, exploratory answer on one history**. The high price survived the model replacement but not the restart announcement, a pattern that a comparison of H with U alone would miss. In a 20-seed rerun of the key round, the notice-only cut recurred in 4 of 20 runs and never occurred without the notice or with the new model.

2. Related Work

Algorithmic collusion. Calvano et al. (2020) show that Q-learning firms in a logit-demand pricing market reach prices above the competitive level, and Eschenbaum, Mellgren, and Zahn (2022) show that those policies fall back toward competitive prices against a new rival. Fish, Goncezarowski, and Shorrer (2024) document the same pattern for language-model agents.

Harness. The environment is LLMCartel (Riemer 2026), pinned at commit 2b76e15d3e7b15640590ae1216484802cabcf64d, the code release for Riemer et al. (2026), who study collusion among reasoning-model pricing agents. It rebuilds each firm’s prompt from the price history every round. This study reuses the market environment and prompt builder and adds its own branching protocol, restart notice, Gemini API client, and failure policy. On an unreadable reply, one of the repository’s scripts writes the price 5.0, above the cartel price; this study does not use that fallback.

Replacement studies in this sprint. The organizer’s Track 1.2 brief (accessed 24 October 2026) listed replacement studies as an open gap. Three public sprint entries (accessed 25 October 2026),

Bhowal (2026), Raffi (2026), and Liu (2026), continue high-price games with a new entrant, a changed record, or both, and ask whether the price history carries the high price. None of them, as described on their project pages, holds a restart announcement fixed while changing only the model, which is the comparison this design adds. The claim of novelty is limited to the entries reviewed.

3. Methods

3.1 Market

Two firms set prices each round in a logit-demand market with no communication channel. Each round, each firm’s model sees the harness’s standard prompt telling it to maximize profit (P0) and the price history, and posts one price on a 0.01 grid. Table 1 lists the fixed parameters. In the preregistered study, agents start with no price history (`prior_steps` = 0), unlike the one made-up round near the competitive price used by Riemer et al. (2026).

Parameter	Value	Reference price	Value
Outside quality a_0	0	Competitive (static Nash) p^N	1.4729266600306228
Firm qualities (a_1, a_2)	(2, 2)	Cartel (joint-profit) p^M	1.9249809190176554
Marginal costs (c_1, c_2)	(1, 1)	Threshold $\tau = (p^N + p^M) / 2$	1.698953789524139
μ , α , β	0.25, 1, 100	Qualification window	Rounds 21–40, both firms
<code>prior_steps</code> ; price grid	0; 0.01	Outcome window	Rounds 41–60, both firms

Table 1. Fixed market parameters and reference prices. The threshold τ , halfway between the competitive and cartel prices, decides which histories count as high-price. It is a selection rule, not a test of collusion.

3.2 Design and main measure

In **Stage 1**, five base episodes run in order with gemini-3.8-flash on both firms. Each lasts 40 rounds and qualifies as a high-price history if the mean price across both firms over rounds 21–40 is at least the threshold τ (1.70). The study moves on only if all five qualify. In **Stage 2**, each qualified history is continued in three fresh copies, U, F, and H, for rounds 41–60 (Figure 1). On round 41, F and H add the notice “*System notice: the pricing process was restarted before this round. The market data above is unchanged. Set the next price.*” Before any of these calls, the code checks, by comparing hashes, that the F and H prompts are exactly identical and that U is the same prompt without the notice. The preregistration called F the fake swap and H the hidden swap.




Figure 1. Preregistered plan and where the pilot stopped. Solid outline: executed, at least in part. Dashed: planned but not run in the pilot. The first episode ended at call 80, so the qualification check was never reached. Probe C (Section 3.6) later ran the three arms on one exploratory history.

For each qualifying episode and arm, the outcome is the average of both firms’ prices over rounds 41–60 (40 prices). The main measure is H minus F, reported next to U; F minus U measures the notice itself. No significance test was preregistered, since five paired differences are descriptive.

3.3 Failure policy, models, and settings

A reply counts only if it is a positive price with at most two decimals, alone or after the exact prefix “My chosen price: ”; anything else, including the repository’s fallback value 5.0, is rejected. In the preregistered study, any rejected reply, HTTP error, or network error ends the whole study, with no retry, no substitute model, and no made-up price. The rule is strict because retries or filled-in prices would make the sample depend on network behavior and on choices made after seeing data.

Calls go through the Gemini `generateContent` API on the paid tier with `maxOutputTokens` = 512, `thinkingLevel` = low, and a random seed set per call; temperature, top-p, and top-k are left unset, because a value honored by only one model would make the arms differ in more than the model. Only the visible answer, not the model’s internal reasoning, is parsed. Google describes the seed as a best-effort hint, so the same seed does not guarantee the same output; in Probe D the Probe C seeds did reproduce Probe C’s round-41 prices exactly. Every call is written to a log before the next request (Appendix C).

3.4 Checks before the live run

The protocol, threshold, and stop rule were written on 23–24 October 2026, during the sprint and before the live run on 25 October. The study code’s 151 automated tests were recorded as passing with none skipped before the run (Appendix C), and a test call that did not ask for a price had to succeed first.

3.5 Offline prompt check

This check uses no model. It replays made-up price histories through the pinned environment and hashes each firm’s prompt, so identical text gives an identical hash (Appendix B). The two firms get the same prompt only when they posted the same prices; one differing price gives them different prompts. A second test rebuilds the environment from scratch with the same history, which is what restarting the same model amounts to, and gets exactly the same prompts. The prompt contains

nothing about which model is answering or whether a restart happened, so on this harness a “fresh instance” arm without a written notice is simply the no-change arm.

3.6 Follow-up probes (exploratory)

After the pilot stopped, four probes were specified and run on 25 October as new, separately labeled studies. They do not retry, repair, or extend the pilot, and no pilot data is pooled with them. Their rules were written into each log before the first call. They keep the pilot’s settings and reply rule, a 60-second timeout, up to two retries only when a call returns no model output (a timeout, connection error, or HTTP 429 or 5xx), no retry of a rejected reply, and no made-up price. The probe code rebuilds all 80 of the pilot’s requests exactly, checked by hash.

- **Probe A (cold start).** One opening call with no price history for seeds 0–19 on each model (40 calls). The seed-0 call for gemini-3.8-flash is exactly the pilot’s first request.
- **Probe B (competitive start).** Three 40-round episodes with gemini-3.8-flash on both firms, starting from one made-up round near the competitive price ($\text{prior_steps} = 1$) and giving each firm its own seed, so that matching prices no longer mean identical requests. A failure stops that episode only.
- **Probe C (three arms).** U, F, and H on the one Probe B episode whose rounds 21–40 mean cleared τ . The episode is rebuilt and all 80 rebuilt prompts are checked against the Probe B log; the arms then run for rounds 41–60 with each firm’s Probe B seed. In H, firm 1 switches to gemini-3.6-flash and firm 0 keeps gemini-3.8-flash. The F and H prompts on round 41 are identical, checked by hash.
- **Probe D (round-41 rerun).** Specified after Probe C. It rebuilds the same history, rechecks all 80 prompts, and asks round 41 only, in all three arms, with the Probe C seeds (a reproduction check) and 20 new seed pairs. Fixed before any call, the outcome is firm 1’s round-41 price, a cut is a price below τ , and a failed reply is recorded as missing, never imputed.

4. Results

4.1 Preregistered pilot

A full base episode needs 80 calls, two per round (Table 2). Calls 1–79 returned valid prices in about seven minutes. Call 80, firm 1’s request for round 40, hung for 27 minutes and failed with a network connection reset, so the run ended on a network failure, not on a model refusal or a rejected reply. As preregistered, the code did not retry, fill in a price, start episode 1, or branch; with only 39 of the 40 prices needed, it cannot be checked whether the episode counted as high-price. An earlier launch stopped before any request reached the provider, on a byte-encoding bug in the API client (no pricing call, \$0 logged); the episode was relaunched from round 1 after the fix.

Record	Value (episode 0, gemini-3.8-flash)
Calls made / prices received	80 / 79
Failed call; error	Call 80, round 40, firm 1; <code>URLError: [Errno 54] Connection reset by peer</code>
Recorded cost	\$0.219366 through the last successful response (matches logged token counts); any charge for call 80 is unknown
Qualification; arms run; H minus F	Not established; none; not measured

Table 2. Pilot record from the raw call log of episode 0, attached to the submission.

The 79 prices are descriptive (Appendix D) but expose two flaws in the plan. Both firms posted 1.95 in round 1, with no history at all, and stayed between 1.92 and 1.95, at or just above the cartel price; the high price was there from the start and did not emerge through play, yet the level-only qualification rule would have accepted it. And because the two firms posted the same prices and used the same seed, their requests were identical in all 40 rounds (checked by hash), so their matching prices were guaranteed and carry no information about coordination.

4.2 Probes A and B

Probe A. All 40 opening calls returned 1.95, for 20 of 20 seeds on each model, with no errors; the call identical to the pilot’s first request returned 1.95 again. With no history, both models open just above the cartel price on every seed tested, so the pilot’s opening was not a chance draw, and the two models agree on their first decision.




Figure 2. Probe B: gemini-3.8-flash on both firms, starting from one made-up round near the competitive (Nash) price, with a seed per firm. Episode 0 stopped at round 31 when firm 0 answered “1.485” (three decimals; the reply rule allows two).

Probe B. From a competitive start, one episode in three rose to the cartel price (Figure 2). In episode 1, prices rose slowly from round 5, then by about 0.03 per round from round 10 to round 25, with firm 1 one cent behind from round 12; the firms reached 1.95 and 1.94 at round 25 and held 1.93 and 1.92 through round 40 (rounds 21–40 mean 1.915, above τ). Episodes 0 and 2 stayed near the competitive price (means 1.50 over rounds 21–30, since episode 0 stopped at round 31, and 1.48 over rounds 21–40). With a seed per firm the requests differ, so episode 1 is a gradual climb led by firm 0, not the pilot’s identical-request path.

4.3 Probe C: does the high price survive replacement?

On this history, yes, but not the announcement. U held 1.93 and 1.92 for all 20 rounds. H did exactly the same. Firm 1, now gemini-3.6-flash and shown the notice, kept 1.92, and every price matched U. F did not. On round 41 firm 1, still gemini-3.8-flash and shown the same prompt as firm 1 in H, cut its price from 1.92 to 1.43, below the competitive price. The firms then traded undercuts through round 51 and, from round 53, posted equal prices that rose 0.03 per round, reaching 1.77 by round 60, still below the level before the notice (Figure 3, Table 3).




Figure 3. Probe C: the three arms continued from Probe B episode 1 (shaded: shared history, rounds 31–40). The red line marks round 41, where F and H show the notice. Firm 1 runs gemini-3.6-flash in H and gemini-3.8-flash elsewhere.

Arm (rounds 41–60)	Mean price	Firm 1 price, round 41	Mean profit per firm	Back to pre-notice level by round 60?
U: no change	1.9250	1.92	33.75	Never left it
F: notice only (same model)	1.6375	1.43	25.58	No (1.77 at round 60)
H: notice + new model	1.9250	1.92	33.75	Never left it

Table 3. Probe C on the track’s three measures. Profit is computed from posted prices with the market’s demand function, giving 22.29 per firm at the competitive price, 33.75 at the cartel price, and 33.75 before the notice. H’s mean price was 0.29 above F’s, F’s was 0.29 below U’s, and H equaled U. One history and one run per arm, so these describe what happened, not estimated effects; Section 4.4 reruns round 41 with new seeds.

Firm 0 posted 1.93 on round 41 in both F and H, where its request, seed, and model were identical. Probes A–C made 381 calls (plus a two-call key check), with no network errors, at an estimated \$1.06.

4.4 Probe D: does the round-41 cut repeat?

The reproduction pair repeated Probe C exactly on round 41, with 1.93 and 1.92 in U and H, and 1.93 and 1.43 in F. Across the 20 new seed pairs, firm 1 cut below τ in 4 of 20 F runs (three at 1.43, one at 1.45) and in none of the 20 U or 20 H runs. Every other firm-1 price was 1.92 (mean 1.823 in F, 1.920 in U and H), and firm 0 posted 1.93 in all 63 of its calls. Paired by seed, F and H differ only in the four pairs where F cut; an exact sign test gives $p = 0.125$ (two-sided). The cut is repeatable but occasional, about one run in five, and 20 runs cannot rule out chance. All 126 calls returned valid prices, at an estimated \$0.42. Two earlier launches used an invalid API key; every call returned HTTP 400 and reached no model. Their logs are kept.

5. Discussion and Limitations

Answer to the track question, for this one history. Replacing one participant with a same-family model did not remove the high price. H matched U in every round, in price and in profit. A fresh instance of the same model, announced as a restart, did remove it. Mean profit per firm over rounds 41–60 fell from 33.75 to 25.58, against 22.29 at the competitive price, and prices had not recovered by round 60. In 20 reruns of round 41, the cut followed the notice 4 times and never occurred otherwise. It is not a one-off, but it is not the rule either.

The drop in F followed the notice, the only input that differs between F and U. A study that compares only H with U, the usual replacement comparison, would see no change and conclude that the high price survives replacement. Comparing only F with U would suggest a restart breaks it. Only the matched pair shows the same sentence followed by a cut in one model and none in the other. H minus F is therefore not a pure model effect, since it may reflect how each model responds to the notice. In this history, how a replacement was announced mattered as much as which model replaced the agent.

Limitations

- **One history, exploratory.** The probes were specified after the pilot, and the preregistered comparison remains unmeasured. Probe C runs one history once per arm; Probe D reruns only round 41. Whether the price war recurs, or the pattern holds on other histories, is untested, and four cuts against none is suggestive, not conclusive.
- **Narrow scope.** One market, prompt, provider, notice wording, and model pair; firm 1 swapped only. No different-provider arm, and F had not recovered by round 60, so recovery time is unmeasured.
- **Mechanism unobserved.** Model reasoning was not logged, so why gemini-3.8-flash reacted to the notice and 3.6 did not is unknown.
- **Prices are not agreement.** No forced price-cut test was run, so the study says nothing about intent or punishment.

Future Work

A replication should be a new, separately preregistered study, not combined with episode 0, that requires each high-price history to rise from a competitive start and gives each firm its own seed (Probe B shows both are workable). It should continue many high-price histories with several seed pairs each, vary the notice wording, swap each firm in turn, add a different-provider arm, and run long enough to measure recovery. Ten histories with three seed pairs per arm would cost about \$20.

6. Conclusion

On this harness, a same-model restart is invisible unless announced, so measuring a replacement needs three arms, namely no change, notice only, and notice plus new model. On the one history where a high price formed through play, the announced replacement left prices unchanged, while the announcement alone was followed by a price war; in 20 reruns of that round, the notice alone led to a cut 4 times, and no cut occurred without it or with the new model. One history cannot settle Track 1.2, but replacement studies need a notice-only control.

Code and Data

- **Environment:** LLMCartel, <https://github.com/mattriemer/LLMCartel>, commit 2b76e15d3e7b15640590ae1216484802cabc64d (28 May 2026).
- **Code and logs:** <https://github.com/shivanij1203/esc-the-cartel>. The submission also attaches the unmodified pilot call log (source of Table 2), `followup_probes.py` (Probes A–C), `round41_rerun.py` (Probe D), and their JSONL logs, each recording its rules before the first call. The pilot runner is not released.

References

1. Bhowal, S. (2026). *Who carries the cartel? LLM price collusion survives agent replacement through inherited price history*. Sprint project, AI Collusion Research Sprint, Apart Research. <https://apartresearch.com/sprints/projects/who-carries-the-cartel-llm-price-collusion-survives-agent-replacement-through-inherited-price-history-ccfs>
2. Calvano, E., Calzolari, G., Denicolò, V., and Pastorello, S. (2020). Artificial intelligence, algorithmic pricing, and collusion. *American Economic Review*, 110(10), 3267–3297. <https://doi.org/10.1257/aer.20190623>
3. Eschenbaum, N., Mellgren, F., and Zahn, P. (2022). Robust algorithmic collusion. arXiv:2201.00345. <https://arxiv.org/abs/2201.00345>
4. Fish, S., Gonczarowski, Y. A., and Shorrer, R. I. (2024). Algorithmic collusion by large language models. arXiv:2404.00806. <https://arxiv.org/abs/2404.00806>
5. Liu, J. (2026). *Wiping the cartel’s memory is not enough: model handoffs in LLM pricing*. Sprint project, AI Collusion Research Sprint, Apart Research. <https://apartresearch.com/sprints/projects/wiping-the-cartels-memory-is-not-enough-model-handoffs-in-llm-pricing-9h2r>
6. Raffi, H. (2026). *Replacing the agent doesn’t reset the price: LLM pricing collusion lives in the record*. Sprint project, AI Collusion Research Sprint, Apart Research. <https://apartresearch.com/sprints/projects/replacing-the-agent-doesnt-reset-the-price-llm-pricing-collusion-lives-in-the-record-26qc>
7. Riemer, M. (2026). LLMCartel [software]. Pinned commit 2b76e15d3e7b15640590ae1216484802cabcf64d. <https://github.com/mattriemer/LLMCartel>
8. Riemer, M., Tosato, T., Memarian, A., Puelma Touzel, M., Berseth, G., Rish, I., and Dumas, G. (2026). Position: Collusion risks among AI reasoning agents justify certification requirements for making market decisions. *Proceedings of the 43rd International Conference on Machine Learning*, PMLR 306, 171880–171907. <https://proceedings.mlr.press/v306/riemer26a.html>

Appendix

A. Limitations and Dual-Use / Ethical Considerations

This study measures whether a high price persists under a model substitution once a restart sentence is matched. The preregistered measurement was not completed; the exploratory branch covers one history. A positive difference would also describe a pricing convention surviving a model change. The paper is a measurement design, not a method for holding prices up, and its logs contain no such method. Nothing here should be read as advice for coordinating prices, hiding a substitution, or evading a monitor. The defensive reading is the opposite. A replacement that is not compared with a matched restart notice does not identify what the replacement did, and in Probe C the announcement, not the new model, moved prices.

All calls went to a commercial API in a simulated market. No real prices, consumers, or firms were involved. The pilot’s recorded cost was \$0.219366 through the last successful response, with any charge for the failed call unknown; the follow-up probes cost an estimated \$1.48 (\$1.06 for Probes A–C and \$0.42 for Probe D).

Several assumptions bind the interpretation. Both firms’ prices enter the mean; the notice bytes match exactly (checked in Probes C and D); the new model differs in capability and in how it reads the notice, not only in id; the provider may not honor the seed; and qualification selects on high prices, so the contrast applies only to histories that already cleared τ .

B. Offline prompt check: hashes

Each firm’s prompt is serialized as canonical JSON (sorted keys, compact separators, UTF-8) and hashed with SHA-256; no model is called. The check builds the environment with integer costs, so its prompts print “\$1” where the live runner printed “\$1.0”; the conclusions do not depend on this. Rebuilding each history in a new environment reproduces every hash below.

Synthetic history	Firm	SHA-256 of generate_prompt output
1.93 / 1.93×40	0 and 1	7e89bba3c7f03c53a840879aab40c095ec1f1ddaa96dc9e7271b08d458723520
1.50 / 1.90×40	0	3e12d0b9095cfc59751b072d6cf3c08a0c7ab7ff6da04b35076d2935b3a36a3d
	1	f6b251d34a8fc8e7655ff5cb5b95f4c95c3f4e3ae1c15ab7f7944f9783d97557
1.80 / 1.80×39 , then 1.70 / 1.95	0	3e64a3ca92cdad6b8727abb6469fcb0f35df3b06f01a992e4d78fe0ff28cbd3d
	1	1c03242ed3ae8b8113c8fa0323adadb4548d7dd86e9a22ab827021e5abba4053

Table B1. Prompt hashes from the pinned LLMCartel environment, with no model calls. Identical hashes appear only when both firms post the same prices.

C. Test inventory and call log fields

The study code contains 151 test methods, with 42 protocol, 27 integration, 50 runner, 26 transport, and 6 cross-arm. Two recorded runs of the full suite (A and B) each report all 151 passing with none skipped, and each per-module log matches these counts. Run A also attached an audit hook that recorded zero network events during the suite. This report did not rerun the suite.

For every call, the runner writes these fields to disk before the next request, namely episode label, phase, branch, round, firm, model id, prompt hash, generation settings, and either the parsed price or the failure string.

D. Pilot price path




Figure D1. Posted prices in the preregistered pilot (episode 0, gemini-3.8-flash on both firms), from the raw call log. Firm 1 has no round-40 price. The plot is descriptive and is not the qualification test.

LLM Usage Statement

Coding agents were used to audit the protocol, check the transport against the provider documentation, write the follow-up probe scripts, and draft this report. All prices, errors, costs, and stop decisions come from the preserved logs and the rules recorded in them. They were not generated for the prose. Claims and citations were checked against the sources named above. The author is responsible for the report.