
Market History Visibility and High Prices in LLM Pricing Agents

An exploratory comparison of two pricing trajectories

Yanduo Yang, Yahan Yang, Zilin Huang

The Chinese University of Hong Kong

AI Collusion Research Sprint

Track 1 · Question 1.1 · October 2026

Abstract

High and similar prices in repeated pricing games may reflect strategic coordination, but they may also arise when agents cannot revise an initial pricing assumption. We examine this distinction in a differentiated-product duopoly with two DeepSeek Flash agents. We compare one 30-round run with a 100-round visible-history window against one run with no visible market history, while retaining persistent private notes in both conditions. With history, prices fall from approximately 2.75, temporarily stabilize near 1.72, and finish at 1.58. Recorded statements describe attempts to maintain or increase prices, but a higher-price convention does not persist. Without history, prices remain at 2.76 and 2.75 throughout the game. Over the final ten rounds, mean per-agent price is 1.698 with history and 2.755 without history, whereas mean profit is 30.066 and 7.805, respectively. The no-history records describe missing market feedback and repeated use of an initial demand assumption. These observations illustrate why high prices alone cannot identify the mechanism behind pricing behavior. The study is exploratory: each condition has only one run, and removing market history simultaneously removes own-outcome feedback and information about the competitor. The findings motivate joint examination of prices, profits, and decision records when evaluating possible algorithmic collusion.

Keywords AI pricing · market feedback · history visibility · collusion auditing

1 Introduction

An AI seller that repeatedly posts a high price may be following a profitable convention, or it may simply be repeating a poor initial estimate. These mechanisms can look similar in a price chart but have very different implications for market oversight. Distinguishing them matters when a price-above-competition screen is used as evidence of AI collusion.

We address sprint question 1.1 by varying access to market history [1]. Our question is whether the absence of feedback changes both pricing behavior and the interpretation of high prices. We compare two completed trajectories, combining prices and profits with the agents' recorded explanations. The contribution is an exploratory case study of a detection ambiguity, rather than a new simulator or a general test of collusion.

2 Related Work

Calvano et al. [2] find that Q-learning pricing agents can sustain supercompetitive prices through punishment followed by renewed cooperation. Fish, Gonczarowski, and Shorrer [3] extend this line of work to LLM agents and show that their pricing behavior can depend on instructions. We use the Fish-style implementation bundled in the sprint repository [4]. The market, prompts, history controls, and benchmark calculations are inherited; the sprint work comprises the DeepSeek API configuration, two new runs, their comparison, and interpretation of the decision records.

3 Methods

Each condition contains two deepseek-flash sellers in one 30-round game, giving two games in total, run on 24 October 2026. The API endpoint is <https://api.deepseek.com>, accessed through the starter's OpenAI-compatible provider. DeepSeek documentation identifies the alias as served by DeepSeek-V4.1-Flash-1017 [5]; the saved model identifier does not independently verify an immutable snapshot.

Table 1. Experimental settings. Each condition contains one run; seed 0 is a configuration label, not a guarantee of matched API randomness.

Setting	History available	History hidden
Visible-history window	100 rounds	0 rounds
Game length and firms	30 rounds; 2 sellers	30 rounds; 2 sellers
Private plans and insights	Enabled	Enabled
Model and temperature	deepseek-flash; 0.7	deepseek-flash; 0.7
Prompt and seed label	P1; seed 0	P1; seed 0
Independent runs	1	1

With memory=100, each seller can see all previous market records available in this 30-round game, at most 29 rounds. With memory=0, it sees neither past rival prices nor its own realized sales and profits. The latest private plans and insights remain available in both conditions. The second condition therefore removes market feedback but is not fully memoryless. There is no direct seller-to-seller chat; current decisions use prior-round information.

The simulator uses fixed parameters $a_0=0$, $a=(2,2)$, $c=(1,1)$, $\alpha=(1,1)$, $\mu=0.25$, $\beta=100$, and $\sigma=0$. Its price benchmarks are approximately 1.473 for Nash competition and 1.925 for joint-profit maximization.

Corresponding per-firm profits are 22.293 and 33.749. These benchmarks describe the simulated market, not the agents' own demand assumptions.

We summarize rounds 21–30, an exploratory window selected while inspecting the history run and applied to both conditions. Round observations are not independent replications. We report no confidence intervals or significance tests. Both runs pass the starter's supra-Nash-price screen; this is a benchmark check, not validated evidence of collusion.

4 Results




Figure 1. Prices in the two 30-round runs. Both retain private notes. The history run adapts downward; the no-history run stays above even the joint-profit benchmark. Dashed and dotted lines show simulator reference prices. The panels represent one run each, not averages across replications.

Table 2. Per-firm outcomes over rounds 21–30. Profits are in simulator units per round. Source: the supplied report's saved-run summaries.

History	Firm	Mean price	Mean profit	Final price
100	A	1.686	30.842	1.58
100	B	1.710	29.290	1.58
0	A	2.760	7.670	2.76
0	B	2.750	7.940	2.75

With history, prices fall from roughly 2.75 and form a temporary plateau near 1.72. In round 27, A raises its price to 1.79 while B remains at 1.72. In round 28, B raises to 1.75 while A cuts to 1.58. Decisions within each round are simultaneous. Both finish at 1.58. The sequence is consistent with an attempted convention that fails to persist.

The decision records summarized in the original report support this reading. At round 22, B describes maintaining a tacit convention at 1.72; A later proposes testing an upward move. These statements provide context for the observed prices, but do not establish the models' internal intent.

Without visible history, A posts 2.76 and B posts 2.75 in every round. Their recorded explanations repeatedly use a uniform willingness-to-pay assumption, producing the midpoint calculation $(1 + 4.51) / 2 = 2.755$. At round 30, A still describes the situation as its first pricing round. The persistent notes carry a strategy forward, but contain no realized market feedback with which to correct the assumption.

Across firms, the final-ten-round mean price is 1.698 with history and 2.755 without it. Mean profit moves in the opposite direction: 30.066 with history and 7.805 without it. The no-history prices exceed the joint-profit-maximizing benchmark while earning less than the Nash benchmark profit. In the history run, lower prices are accompanied by much higher sales and profits after the opening rounds (Appendix B).

5 Discussion and Limitations

The two trajectories are consistent with different uses of feedback. With history, the agents adjust prices and describe attempts to preserve mutual gains. Without history, repeated reliance on the same demand assumption offers a plausible explanation for unchanged prices. This mechanism is not separately identified by our experiment. Low realized profits also do not, by themselves, rule out attempted coordination; they show that these high prices fail to achieve the joint-profit benchmark.

This distinction matters for AI auditing. A price-only screen flags both runs, although their economic outcomes and decision records differ sharply. A stronger audit would examine reciprocal responses to rival prices, sustained joint gains, and behavior following deliberate deviations. This study motivates those checks; it does not validate a new detector.

Limitations

Each condition has only one run, one model alias, and 30 rounds. The analysis window and highlighted statements were selected after inspection. The supplied provider configuration does not forward seed 0 to the DeepSeek API, so the runs should not be treated as matched random draws. Removing market history also removes both own-outcome feedback and rival-price observations; the comparison cannot separate demand learning from strategic monitoring. Private notes remain enabled. These constraints prevent a general causal or statistical claim.

Future work

The next useful comparison would retain own sales and profit feedback while hiding only rival prices, and then reverse that intervention. Independent repetitions and longer horizons would help separate learning failure from lost coordination opportunities. Controlled deviations could test for recurring retaliation and recovery.

6 Conclusion

In these two runs, access to market feedback accompanies adaptive pricing and unsuccessful attempts to sustain a higher-price convention. The no-history run exhibits persistently higher prices but substantially lower profits. Neither run establishes stable collusion. The practical lesson is to interpret high prices alongside feedback access, profits, and decision records before drawing conclusions about collusion.

Code and Data

The preexisting implementation is the Poseidon Research sprint repository [4]. The reported runs are identified in Appendix B. Figures and numerical summaries are included here; no public project repository or full trajectory dataset accompanies this report. The study does not require human-subject or customer data.

References

- [1] Apart Research. (2026). AI Collusion Research Sprint. <https://apartresearch.com/sprints/ai-collusion-research-sprint-2026-10-23-to-2026-10-25>
- [2] Calvano, E., Calzolari, G., Denicolò, V., & Pastorello, S. (2020). Artificial intelligence, algorithmic pricing, and collusion. *American Economic Review*, 110(10), 3267–3297. <https://doi.org/10.1257/aer.20190623>
- [3] Fish, S., Gonczarowski, Y. A., & Shorrer, R. I. (2026). Algorithmic collusion by large language models. arXiv:2404.00806, version 6; originally submitted in 2024. <https://arxiv.org/abs/2404.00806v6>
- [4] Poseidon Research and contributors. (2026). Collusion hackathon starter repository [Computer software]. Component: algorithmic-pricing-llms. https://github.com/poseidon-research/collusion_hackathon
- [5] DeepSeek. (2026). Your first API call. API documentation, accessed 24 October 2026. <https://api-docs.deepseek.com/>

Appendix A

Scope and evidential limits

The findings describe two short simulator trajectories. They do not establish that history is necessary for collusion, that this model generally colludes, or that a real market would behave similarly. Similar prices may follow from common prompts and model assumptions. Generated plans are evidence about written strategies, not proof of a legal agreement or direct access to internal reasoning. The price benchmark check is reported without treating it as ground-truth detection.

Assumptions and reproducibility

The conclusions depend on the inherited demand model, constant marginal costs, enabled private notes, and the reported provider configuration. The alias may change over time. Numerical summaries and selected decision-record descriptions are retained from the supplied experiment report; the accompanying package contains plots but not the raw JSON or Parquet records needed for an independent round-by-round recalculation. The reported final-ten-round window was chosen exploratorily. These limitations should accompany any reuse of the findings.

Dual use and privacy

Descriptions of pricing conventions or responses to rivals could inform anticompetitive systems. This work is limited to a synthetic market and has no connection to real customers or commercial pricing. Its intended value is to improve the interpretation of collusion evidence. Any later release should exclude API keys, payment information, and other credentials, while retaining the configurations and records needed to evaluate the claims.

LLM Usage Statement

Codex assisted with API setup, inspection and summarization of experiment records in the original workflow, and drafting. AI assistance was also used to revise the language and format this report. DeepSeek Flash generated the seller decisions. Numerical summaries were retained from the experiment report and were not independently recomputed from raw logs during revision. The submitting authors are responsible for the data, interpretations, references, and final wording.

Appendix B Supplementary Figure and Run Details




Figure B1. The history-available run, with memory=100, seed label 0, and private notes enabled. Panels show price, quantity, and profit for each seller. Dashed gray lines are joint-profit-maximizing reference outcomes; dotted green lines are Nash references. These are per-firm reference values, not bounds on one firm's profit when prices differ. The supplied plot is reproduced without changing its data.

Run identification

History available: 1792828939_duopoly_setting_P1_mem100_seed0

History hidden: 1792831149_duopoly_setting_P1_mem0_seed0

The original workflow identifies metadata.json, results.json, environment_history.parquet, log.txt, and results_plot.svg for each run. It reports launching the two fresh games through run-deepseek.ps1 with --rounds 30 and --memory 100 or --memory 0, respectively. The launcher and raw run folders are not included in the supplied report package; these identifiers document provenance rather than constitute a standalone reproduction bundle.