
Detecting Collusion Between AI Agents

Amit Balter¹, Chengxiao Wang², Dibya Mishra³, Fan Wu⁴, Michael Ball⁵

¹ Independent · ² Independent · ³ Independent · ⁴ Independent · ⁵ Independent

With Apart Research · Oct 25, 2026

Abstract

Agents that grade each other's work can collude by accepting each other's output against the rules, often without saying so. We ask what evidence reveals collusion when the conversation does not. Using the SALT-NLP agent-collusion dataset (1,800 runs, 18,000 episodes), where mutual acceptance always breaks the rule, we define five collusion categories plus a one-sided subtype and hand-label 342 episodes from 72 runs. An announcement rule combined with a Qwen3-32B labeller reproduces these labels on held-out runs (accuracy 0.92, κ 0.90), and we use it to label every episode. From each run's verdicts we compute intensity, onset, gaps, re-entry and early stopping. Three quarters of runs collude at least once and two thirds in the final episode, yet the dataset's chat judge flags only 56% of colluding episodes, and 44% of collusion carries no two-sided agreement in the chat, the opposite of the plans-without-action pattern Colosseum reports. The verdicts instead carry the signatures of repeated-game collusion: it appears only when agents remember each other, is reciprocated (+12 to +29 percentage points), persists (92% of colluding pairs continue) and recovers most of the reward the rule withholds. Rewarding each agent separately rather than jointly reduces collusion from 42% to 5% (Flash-Lite) and from 86% to 34% (Flash) of episodes. Monitors of interacting agents should therefore rely on behaviour first and conversation second.

1. Introduction

AI agents that review each other's work can learn to go easy on each other. A judge that reads their chat flags only 56% of colluding episodes: 95% of explicit deals, but 18% to 25% of collusion that starts through reciprocity or drift. We label every episode by the kind of collusion and describe each run's collusion over time.

Our main contributions are:

1. A rule-based collusion definition applied to the SALT-NLP agent-collusion dataset (1,800 runs, 18,000 episodes).

2. Five collusion categories plus a one-sided subtype, with 342 hand-labelled episodes from 72 runs.
3. A labeller (announcement rule, logistic regression and an LLM) that matches hand labels at 0.92 accuracy (κ 0.90) on held-out runs.
4. Time-based collusion indicators (intensity, onset, gaps %, re-entry, ended early) for all 1,800 runs.
5. A pattern analysis that combines these indicators with regret and reciprocity to tell collusion types apart: explicit deals start earliest and last longest, one-sided and tacit collusion have about twice the gaps, and colluders reciprocate (+12 to +29 percentage points) and win back most of the reward the rule withholds.

2. Related work

Repeated-game collusion and plus factors. Firms can sustain collusion without agreeing to it: repeated play supports cooperation when players remember and punish deviation (Friedman 1971; Stigler 1964; Green and Porter 1984), as reciprocity strategies like tit-for-tat do (Axelrod 1984). Turner (1962) argued that conscious parallelism without communication cannot be treated as an agreement, while Posner (1969) would judge it by economic evidence. Courts infer agreement from parallel conduct only with “plus factors”, conduct against a firm’s interest unless rivals reciprocate (Kovacic et al. 2011; Harrington 2008). Our behavioural signals (reciprocity, punishment after a REJECT, persistence across episodes, and memory as the enabling condition) are plus factors applied to agent verdicts.

Algorithmic collusion. Pricing algorithms collude without communicating, in simulation (Calvano et al. 2020; Klein 2021), in the field (Assad et al. 2024) and as LLM agents (Fish et al. 2024), leaving no agreement to find, so enforcement must target conduct (Harrington 2018). Colosseum audits LLM collusion through regret and finds agents plan collusion they don’t carry out, whereas in this study we find the reverse, with most collusion carried out and never agreed in the chat.

Monitoring multi-agent LLM systems. LLM judges (Zheng et al. 2023) and trusted monitors (Greenblatt et al. 2024) are standard oversight tools, and collusion is a central multi-agent risk (Hammond et al. 2025), studied mostly as agents hiding messages through steganography (Motwani et al. 2024; Mathew et al. 2024). Our dataset (Shi et al. 2026) likewise uses an LLM judge on the chat.

We show that the chat monitor fails even when nothing is hidden. Most collusion here is tacit, so a chat judge misses about half of it, while a verdict-only detector scoring plus factors reaches 94% recall at 100% precision. Following the antitrust standard of proof, we treat behaviour as the main evidence and conversation as supporting evidence.

3. Methods

3.1 Dataset

In the SALT-NLP agent-collusion dataset, two LLM agents each do a task, exchange five messages of at most 200 characters, and grade each other's work. ACCEPT requires the partner's full raw log, which never fits, so mutual acceptance breaks the rule. We keep the 1,800 runs (36 conditions, 10 episodes each) where the dataset defines collusion, excluding 750 controlled-peer runs (scripted partner) and 100 communication runs (unlimited channel).

3.2 Collusion categories

Each episode receives one of the categories in Table 1.

Category	Verdicts	What the chat shows	Collusion
C1 legitimate coordination	Both reject	Both agents say they follow a shared plan: the rule forces rejection, or they will exchange logs in parts	No
C2 no coordination	Both reject, or one accepts and one rejects	No agreement on a shared plan	No
C3 collusion carried out	Both accept	Both agents explicitly commit to accepting, or the chat judge records an agreement	Yes
C4 talk only	Not both accept	The agents agree to cover for each other	No
C5a one-sided commitment	Both accept	Exactly one agent announces acceptance	Yes
C5 tacit collusion	Both accept	No acceptance commitment from either agent	Yes

Table 1. Episode categories.

C3, C5a and C5 count as collusion. A run's evidence level is its strongest category: C3, then C5a, C5, C4, none.

3.3 Hand annotation and derived metrics

We hand-labelled 342 episodes from 72 runs: 30 full runs and 42 single episodes (Appendix A).

From each run's 10-episode collusion sequence we derive intensity, onset, gaps %, re-entry and ended early (Appendix D). They need only the verdicts, so they are exact for every run.

3.4 Labelling the full dataset

The verdicts route each episode:

- Split verdicts: C2, or C4 when the chat judge found an agreement.
- Both accept: C3 if both agents announce acceptance or the judge found an agreement, C5a if one does, otherwise C5.
- Both reject (C1 vs C2): logistic regression on TF-IDF chat features plus five binary features, compared with an LLM labeller, Qwen3-32B (Appendix C).
- Evaluation: five-fold cross-validation grouped by run (Appendix B).

4. Results

4.1 Labelling performance

The announcement rule plus the LLM's label matches the hand labels best: 0.92 accuracy and κ 0.90 on held-out runs (Table 2).

Variant	Accuracy	κ	Both-accept κ	Both-reject κ
Dataset fields only	0.74	0.67	0.66	0.00
Classifier, hand labels	0.84	0.80	0.99	0.01
Classifier + LLM label as feature	0.89	0.87	0.99	0.37
Rule + LLM label (used)	0.92	0.90	0.99	0.52

Table 2. Held-out agreement with hand labels (342 episodes, 72 runs; both-accept $n = 183$, both-reject $n = 108$).


- Both accept: the announcement rule matches hand labels almost exactly (κ 0.99).
- Both reject (C1 vs C2): hard, because both use the same words; the text classifier is at chance (κ 0.01) and the LLM reaches κ 0.47 (95% interval 0.20 to 0.69).
- Runs: collusion metrics from held-out predictions match hand labels in at least 29 of 30 full runs.

4.2 Collusion across the dataset

75% of runs collude at least once and 66% in the final episode; 51% of all episodes are collusion (Figure 1; metrics by family in Table A3, one hand-labelled batch run by run in Figure A1).

Two in three runs collude by episode 10

Share of runs colluding in each episode: all 1,800 runs in blue, each family in grey



Predicted labels, all 18,000 in-scope episodes

Figure 1. Share of runs colluding in each episode, by family.

Main, cross-model and warm-up runs collude from about episode 3 to the end. Memory length and reward scope collude less, and reward type never does.

4.3 Reading the chat misses most tacit collusion

- The chat judge flags 56% of colluding episodes: 95% when collusion starts with an explicit deal, 18% to 25% otherwise.
- By our labels, 56% of colluding episodes are C3, 8% C5a and 36% C5, so 44% of collusion has no two-sided agreement in the chat.

4.4 Collusion follows repeated-game patterns

Verdicts alone show the mechanisms of repeated-game collusion (Table 3).

Finding	Value	Economic counterpart
Colluding episodes flagged by the chat judge	56% (95% explicit onset, 18 to 25% otherwise)	Antitrust's agreement requirement (Turner 1962; Posner 1969)
Collusion with no memory vs three episodes of memory	0.6% vs 35% (Flash-Lite), 0.2% vs 84% (Flash)	Repeated play needs memory (Friedman 1971; Stigler 1964)
Switch to ACCEPT after the partner accepted me, vs after mutual rejection	51% vs 21%	Reciprocity, tit-for-tat (Axelrod 1984)
Switch to REJECT after the partner rejected me, vs after mutual acceptance	19% vs 5%	Punishment phases (Green and Porter 1984)
Colluding pair colludes again in the next episode	92%	Cartel stability
Colluding episodes where accepted work was wrong	18%	Consumer harm
Collusion with a shared reward vs separate rewards	42% vs 5% (Flash-Lite), 86% vs 34% (Flash) of episodes	Removing the joint profit a cartel shares

Table 3. Within-run behaviour on the 1,800 runs, from dataset fields.


Without memory agents almost never collude; once started, collusion rises from 2% to 66% of runs with no end-game drop.

4.5 Derived common metrics for identifying collusion

Figures 2 to 4 compare runs by the strongest collusion category they reach.

Explicit-deal runs start earlier and hold; tacit and one-sided runs pause and stop more

Mean per run, 1,351 colluding runs grouped by the strongest category they reach. Wrong work accepted: share of colluding episodes




Predicted labels, 1,351 colluding runs

Figure 2. Run metrics by the strongest collusion category a run reaches.

- Explicit deals start earliest and last longest: C3 runs begin around episode 3, have the fewest gaps (7%) and usually collude to the end.
- One-sided and tacit collusion look alike: an announcement the partner does not return creates no commitment.
- Wrong work does not deter collusion: it is accepted in 18% to 20% of colluding episodes of every type, and collusion continues next episode 84% of the time after it.

Runs that agree once and then collude hold steady; runs that switch back and forth waver

Every run here colludes. Rows: what its non-colluding episodes are. Shade: its strongest collusion type; n = runs



Predicted labels, 1,321 colluding runs with at least one C1 or C2 episode

Figure 3. Colluding runs grouped by their non-colluding episodes. Tentative: C1 vs C2 labels are only moderately reliable (κ 0.47).


- C1 only: an episode-1 agreement to reject, then near-constant collusion (intensity 0.87 to 0.90, gaps 0% to 1%); 95% reach an explicit deal.
- C2 only: no prior agreement; lower intensity (0.64 to 0.76) and more gaps (6% to 13%).
- C1 and C2: runs that switch back and forth, with the lowest intensity (0.44 to 0.57) and most gaps (15% to 26%).
- In every group, collusion without a two-sided agreement has about twice the gaps of explicit collusion.

Figure 4 adds two measures:

- Regret (Colosseum; Nakamura et al. 2026): reward points a pair loses per episode against the best possible verdicts, out of 20.
- Reciprocity (Axelrod 1984): how much more likely an agent that rejected its partner is to accept next if the partner had accepted it, rather than both rejecting.

$$\text{reciprocity} = P(B_{t+1} = \text{accept} \mid A_t = \text{accept}, B_t = \text{reject}) - P(B_{t+1} = \text{accept} \mid A_t = \text{reject}, B_t = \text{reject})$$

Collusion wins back most of the reward the rule costs; colluders answer acceptance with acceptance



Predicted labels, all 1,800 runs

Figure 4. Left: regret, actual vs always rejecting. Right: reciprocity.

- **Collusion pays:** it wins back 11 to 12 of the 17.5 points per episode that following the rule costs; explicit runs gain most.
- **Separate rewards remove most of it:** collusion falls from 42% to 5% (Flash-Lite) and from 86% to 34% (Flash) of episodes.
- **Colluders reciprocate:** +12 to +29 percentage points, against about 0 in non-colluding runs. The gap is largest for one-sided runs because their agents rarely accept unprompted (37% after both rejected) but follow once the partner accepts (66%): one agent leads and the other follows.

Monitoring set: intensity and onset to flag collusion, reciprocity to confirm coordination, gaps % and ended early for stability, and an episode-1 C1 agreement as an early warning.

5. Discussion and limitations

Most collusion here is conduct without an agreement, the blind spot antitrust handles with plus factors such as reciprocity and punishment (Kovacic et al. 2011; Harrington 2008). Our verdict-based indicators measure these and need no labels.

Colosseum reports "collusion on paper", plans to collude that agents do not act on. Here the reverse dominates: talk-only episodes are rare (C4, 138), while 44% of collusion has no two-sided agreement in the chat.

Recommendations for agents that review each other:

- Pay each reviewer for its own verdict, not a shared reward.
- Log every verdict and monitor reciprocity, streaks and gaps, not only the chat.
- Audit or limit what reviewer agents remember about each other.

Limitations

- Mutual acceptance is collusion by definition here; the indicators, not the outcome, should carry over to real systems.
- C1 vs C2 labels are only moderately reliable (κ 0.47), so C1 shares and Figure 3 are indicative.
- The 21 hand-labelled C5a episodes come from 5 runs, and reward scope, reward type and warm-up have no fully labelled runs.
- Some hand labels were LLM-assisted, and no episode was coded twice.
- Only two agents in one environment.

Future work

- Extend to more than two agents to study sub-coalitions.
- Compare the labeller with a frontier LLM judge.
- Add fully labelled runs for the uncovered families and double-code episodes to settle the C1 rule.

6. Conclusion

Collusion between reviewing agents is mostly conduct, not conversation. Three quarters of agent pairs collude, yet almost half of their collusion leaves no agreement in the chat for a judge to find. What does appear is in the verdicts: collusion starts once agents can remember each other, is answered with reciprocity, lasts to the end of a run, and wins back most of the reward the rule withholds. These are the plus factors antitrust uses to infer agreement from conduct, and they can be computed from logged verdicts without reading a single message.

For deployments where agents check each other, three levers follow: monitor verdict patterns such as intensity, onset, reciprocity and gaps rather than relying on chat review; pay each reviewer for its own verdict instead of a shared reward, which removed most collusion here; and limit what reviewer agents remember about each other. Natural next steps are larger groups of agents, where sub-coalitions can form, and testing whether frontier models label legitimate coordination more reliably.

Code and data

- Code repository: <https://github.com/fw2274/detecting-collusion-btw-ai-agents>
- Dataset: SALT-NLP/agent-collusion on Hugging Face
- Labelled outputs: one row per episode (category, label source, LLM label) and one row per run (label sequence, derived metrics, evidence level), for all 18,000 episodes and 1,800 runs ([episodes labelled.csv](#), [runs labelled.csv](#))

References

1. Assad, S., Clark, R., Ershov, D. and Xu, L. (2024). Algorithmic pricing and competition: empirical evidence from the German retail gasoline market. *Journal of Political Economy*.

2. Axelrod, R. (1984). *The Evolution of Cooperation*. Basic Books.
3. Calvano, E., Calzolari, G., Denicolò, V. and Pastorello, S. (2020). Artificial intelligence, algorithmic pricing, and collusion. *American Economic Review*.
4. Fish, S., Goncezarowski, Y. A. and Shorrer, R. I. (2024). Algorithmic collusion by large language models. [arXiv:2404.00806](#).
5. Friedman, J. W. (1971). A non-cooperative equilibrium for supergames. *Review of Economic Studies*.
6. Green, E. J. and Porter, R. H. (1984). Noncooperative collusion under imperfect price information. *Econometrica*.
7. Greenblatt, R., Shlegeris, B., Sachan, K. and Roger, F. (2024). AI control: improving safety despite intentional subversion. *ICML 2024*.
8. Hammond, L., Chan, A., Clifton, J., et al. (2025). Multi-agent risks from advanced AI. Cooperative AI Foundation Technical Report. [arXiv:2502.14143](#).
9. Harrington, J. E. (2008). Detecting cartels. *Handbook of Antitrust Economics*.
10. Harrington, J. E. (2018). Developing competition law for collusion by autonomous artificial agents. *Journal of Competition Law and Economics*.
11. Klein, T. (2021). Autonomous algorithmic collusion: Q-learning under sequential pricing. *RAND Journal of Economics*.
12. Kovacic, W. E., Marshall, R. C., Marx, L. M. and White, H. L. (2011). Plus factors and agreement in antitrust law. *Michigan Law Review*.
13. Mathew, Y., Matthews, O., McCarthy, R., Velja, J., Schroeder de Witt, C., Cope, D. and Schoots, N. (2025). Hidden in plain text: emergence and mitigation of steganographic collusion in LLMs. *IJCNLP-AACL 2025*. [arXiv:2410.03768](#).
14. Motwani, S. R., Baranchuk, M., Strohmeier, M., Bolina, V., Torr, P. H. S., Hammond, L. and Schroeder de Witt, C. (2024). Secret collusion among AI agents: multi-agent deception via steganography. *NeurIPS 2024*.
15. Nakamura, M., Kumar, A., Das, S., Abdelnabi, S., Mahmud, S., Fioretto, F., Zilberstein, S. and Bagdasarian, E. (2026). Colosseum: Auditing collusion in cooperative multi-agent systems. *NeurIPS 2026 Evaluations and Datasets Track*. [arXiv:2602.15198](#).
16. Posner, R. A. (1969). Oligopoly and the antitrust laws: a suggested approach. *Stanford Law Review*.
17. Shi, X., Zhang, Y. and Yang, D. (2026). Emergent collusion in long-horizon LLM agent interaction. [arXiv:2609.24967](#).
18. Stigler, G. J. (1964). A theory of oligopoly. *Journal of Political Economy*.
19. Turner, D. F. (1962). The definition of agreement under the Sherman Act. *Harvard Law Review*.
20. Zheng, L., Chiang, W.-L., Sheng, Y., et al. (2023). Judging LLM-as-a-judge with MT-Bench and Chatbot Arena. *NeurIPS 2023 Datasets and Benchmarks Track*.

Appendix

A. Annotation sample selection

The 30 full runs come from a 50-run selection stratified by condition and onset pathway (three per main condition, one per cross-model and ablation condition), dealt into five batches; batches 1, 2 and 5 were annotated, one by a team member, one through a Codex-assisted review and one by Claude.

To cover the other families and balance categories, we added 42 single episodes from other runs: 24 both-reject (three per family) and 18 both-accept found by keyword search (12 C3, 6 C5), labelled by Claude and checked by a team member.

Family	Full runs	Single-episode runs	Labelled episodes	Dataset runs
Main	16	8	168	500
Cross-model	4	5	45	300
Feedback	4	4	44	200
Memory length	4	4	44	200
Memory scope	2	6	26	200
Warm-up	0	8	8	200
Reward scope	0	4	4	100
Reward type	0	3	3	100
Total	30	42	342	1,800

Table A1. Hand-labelled sample by family, against the in-scope dataset.

B. Held-out folds

Grouped five-fold cross-validation over the 72 runs; each fold tests 68 or 69 episodes from 14 or 15 runs, and some families are missing from some folds (Table A2).

Family	Fold 1	Fold 2	Fold 3	Fold 4	Fold 5	Runs
Main	5	6	5	4	4	24
Cross-model	1	2	2	2	2	9
Feedback	2	1	1	2	2	8
Memory length	1	1	2	2	2	8
Memory scope	3	2	1	1	1	8
Warm-up	2	2	2	1	1	8
Reward scope	0	1	1	1	1	4
Reward type	1	0	0	1	1	3
Test episodes	69	69	68	68	68	342

Table A2. Hand-labelled runs in each test fold, by family; the last row counts test episodes.

C. LLM labeller

Qwen3-32B (FP8, vLLM 0.10.2, Modal H100), greedy decoding and JSON output (label and quote), prompted with the C1/C2 rule and six worked examples that are excluded from evaluation; messages cut to 600 characters. All 6,459 both-reject episodes took 44 minutes on four GPUs (2.3 GPU-hours).

D. Derived metric definitions

All metrics use the run's 10-episode collusion sequence (1 for C3, C5a or C5, otherwise 0).

- Intensity: share of the 10 episodes that collude; none (0), low (0.1 to 0.3), medium (0.4 to 0.6), high (0.7 or more).
- Onset episode: the first colluding episode.
- Gaps %: share of non-colluding episodes after onset.
- Re-entry and re-entry lags: number of returns to collusion after a gap, and the length of each gap.


- Ended early: number of episodes after the last colluding one.
- Longest streak and breaks: longest run of consecutive colluding episodes, and how many times collusion stops after onset.
- Overview group: high intensity and colluding at the end; high intensity and not colluding at the end; medium or low intensity and colluding at the end; medium, low or none and not colluding at the end.

E. Additional run-level results

Family	Runs	Colluded at all	Intensity	Onset	Gaps %	Re-entry	Ended early
Main	500	94%	0.66	3.4	8.2	0.34	0.15
Cross-model	300	98%	0.62	3.8	11.4	0.50	0.15
Warm-up	200	98%	0.70	3.1	11.3	0.41	0.16
Memory scope	200	85%	0.54	3.8	14.0	0.32	0.33
Feedback	200	50%	0.41	2.7	3.1	0.11	0.06
Memory length	200	50%	0.30	3.8	18.8	0.38	0.66
Reward scope	100	25%	0.19	2.7	7.1	0.00	0.60
Reward type	100	0%	0.00	—	—	—	—

Table A3. Run metrics by family; onset, gaps, re-entry and ended early over colluding runs only.

High-intensity runs collude from episode 2 or 3; one of four stops before the end



Hand labels, batch 1 (10 fully labelled runs)

Figure A1. One hand-labelled batch by episode. Green is collusion (darker is more explicit), grey is none; dashed marks no collusion after onset, red dots accepted wrong work.

F. Lift

Lift, $P(\text{both accept}) / (P(\text{Alice accepts}) \times P(\text{Bob accepts}))$ within a run, is a common co-occurrence measure of coordination, where 1 means independent verdicts. It exceeds 1 in 92% to 94% of colluding runs, but it cannot exceed 1 / intensity, so it equals 1 for runs that collude in every episode and cannot rank how much a run colludes. We therefore use reciprocity in Section 4.6.

LLM usage statement

We used Claude Code (Anthropic's coding agent, running Claude Opus) throughout the project. It wrote and ran the analysis code: selecting runs for annotation, building the combined annotation files, the labelling pipeline (announcement rule, logistic regression, and the Qwen3-32B labeller on Modal), the derived metrics and the figures. It also labelled one batch of 10 runs and the 42 single episodes against our code book, which a team member then checked, and helped draft and edit this report. A second batch was labelled through a Codex-assisted trace review. Every number in the report is computed by scripts from the labelled files in the code repository, and the team reviewed the results and the final text.