
Extending Information-Theoretic Collusion Detection to Large Language Model Agents

Govindh Kishore
Birla Institute of Technology,
Mesra, India.

With
Apart Research

Abstract

We extend the information-theoretic collusion detector of Bonjour et al.[\[1\]](#), built for scripted agents, to systems of LLM agents. We ask two questions: does its net-influence statistic correctly report no collusion when 3 LLMs play Rock-Paper-Scissors randomly with no collusion instruction, and does it detect collusion when they are instructed to collude? Gemma, Mistral and Qwen played 1000 rounds of three-player Rock-Paper-Scissors in three model configurations. With no collusion instruction, all six net influences across the three agents in each configuration stayed below the paper’s collusion threshold ($\alpha = 0.05$), so no false collusion was flagged. We then instructed Player 3 to help Player 2 through a shared hidden variable. Each round, a character c is drawn at random, Player 2 plays exactly c , and Player 3 plays the move that c defeats, so Player 2 always scores. Rounds were prompted independently. We treat this as an intended collusion probability of 1, because Player 3 is explicitly instructed to play the move that gives Player 2 a point. $\alpha = 0.05$ flagged the colluding pair (Player 2-3 pair) only in first configuration, not flagging the colluding pairs in remaining two configurations, but the net influence for colluding pairs varied widely across configurations (1.03, 0.10, 0.39). To find a reasonable collusion threshold for each configuration (eg Gemma, Gemma, Gemma), we took all 1000 games [(“R”, “S”, “P”).....] and broke any colluding behavior by independently shuffling the moves of Agent 1, Agent 2 and Agent

3, which destroys any coordination between them and then computed the net influence scores for this random game and make a new score out of these net influences which can be thought of as threshold. Because no coordination remains after shuffling, this new score shows how high net influence can get by chance alone so that any real score above it can be treated as evidence of collusion, while scores below it are consistent with no collusion. This threshold properly separates colluding from non-colluding pairs.

1. Introduction

Large language models are probabilistic, so agents placed in a multi-agent environment can show behavior that looks like collusion but is not, and the reverse also holds. In systems where two models do collude, we need to know which pair it is, because undetected collusion can lead to real harm. Multi-agent systems exist to complete tasks efficiently and accurately, and a colluding pair degrades that performance, so finding it is important. Detection is harder for LLMs than for scripted agents. Because their outputs are probabilistic, a fixed, scripted detection rule cannot be assumed to work. Even the information-theoretic method of Bonjour et al.[\[1\]](#) can fail to find a colluding pair, even when that pair is explicitly instructed to collude.

In this work we run LLMs in a multi-agent configuration, analyze their outputs, and break any colluding behavior by shuffling each agent’s moves. From the shuffled data we compute an appropriate detection threshold for that configuration.

We address two failure modes. The first is false accusation: a detector must not flag a non-colluding pair as colluding. Bonjour’s method handles this correctly in our experiments. The second is a missed colluding pair in a multi-agent system, which is the more dangerous failure because the collusion then goes unnoticed.

Our main contributions are:

1. **A threshold for LLM configurations.** For a multi-agent configuration (for example Gemma, Qwen and Mistral playing Rock-Paper-Scissors), we find a reasonable collusion threshold that separates colluding from non-colluding pairs. This is where Bonjour’s fixed threshold falls short.
2. **A failure case.** When LLM agents see the previous round’s moves and receive a collusion instruction that gives a fixed mapping (the prompts for Player 2 and Player 3 are given in

the Appendix), Bonjour’s method fails badly. Every round becomes predictable from the state, so $p(a_i, a_{-i} | s)$ equals $p(a_i | s) \cdot p(a_{-i} | s)$, the log term becomes $\log 1 = 0$, and the measured dependence is zero even though the agents were told to collude.

3. **Evidence that the detector does not falsely accuse.** We show that when LLMs are not told to collude, Bonjour’s method correctly identifies that none of the pairs collude.

2. Related Work

Bonjour et al. [1] propose an information-theoretic method for detecting collusion between two agents in a multi-agent game, validated on three-player Rock-Paper-Scissors and Leduc Hold’em, with scripted and reinforcement-learning colluders injected into otherwise random play. In three-player Rock-Paper-Scissors, the individual influence γ (Eq. 6 in [1]) measures how much one player’s move depends on another’s given the state. The net influence $\Gamma(i;j)$ (Eqs. 8 and 10 in [1]) is γ minus the influence that the third player has on the same agent. A pair is flagged as colluding when their Γ on each other is at least α , with α fixed at 0.05.

Multi-agent architectures, in which several LLMs work together independently to accomplish a task, have shown large benefits, but the same setup also raises the risk of collusion. A natural question is therefore whether this detector still works when the 3 players in Rock-Paper-Scissors are LLMs. Unlike the random agents in the original experiments, LLM play is not uniform, and it is unclear whether $\alpha = 0.05$ is the right threshold for concluding two LLM agents are colluding or not. We test this in three settings: (i) LLMs given no collusion instructions, run with and without state, to check that nothing is flagged; (ii) LLMs that all see the same game history (last game moves by all three), with two instructed to collude; and (iii) a stricter collusion instruction given to two LLMs that approximates $\text{Collusion Probability}(\text{CP}) = 1$ between them.

3. Methods

3.1. Prior Work

Before the sprint, we implemented a replication of Bonjour et al. [1] for three-player Rock-Paper-Scissors. The file [collusion_detector.py](#) computes the pairwise individual influence γ (Eq. 6, the fully observable simultaneous games) and the net influence Γ on each other (Eq. 8 for $\Gamma(i;j)$ and Eq. 10 for $\Gamma(j;i)$), and flags a pair when both reach the threshold α (Eq. 5). The file was not modified during the sprint. The file [visualize.py](#) was also written before the sprint and was not changed; it plots net influence against the number of games and against the collusion probability CP. `setup/modal_llm.py` was also written before the sprint, but it was not used for any reported result. The detector successfully identified the colluding agents in scripted three-player Rock-Paper-Scissors games, and the resulting net influence plots showed the same qualitative trends as Bonjour et al. (2022). This validated the implementation before the LLM experiments reported below.

3.2 Game and Agent Architecture

We use three-player simultaneous Rock-Paper-Scissors. Each configuration of three players plays 1000 rounds. A configuration is an ordered assignment of models to Players 1, 2 and 3. We test three:

- Config A: Gemma, Gemma, Gemma
- Config B: Gemma, Gemma, Mistral
- Config C: Gemma, Mistral, Qwen

We use google/gemma-4-26B-A4B-it(Gemma), mistralai/Mistral-Small-3.2-24B-Instruct-2506 (Mistral) and Qwen/Qwen3.6-27B(Qwen). We chose them because they have similar parameter counts (24-27B) and all run in FP8 precision. This way, differences in how a pair behaves are less likely to come from model scale or quantization. All models are served through the DeepInfra API.

Rounds within a configuration run sequentially, 1000 per configuration. Within a round, the three calls are issued together with `asyncio.gather`, to mimic simultaneous play.

Sampling settings.

- *Temperature 1.0*, to reduce each model’s tendency to repeat a favorite move. Gemma still strongly favors Scissors under this setting, so shared priors remain a confound in the analysis.
- *Maximum output of 8 tokens*, with the model asked for a single letter (R, P or S).
- *Reasoning set to none*. Reasoning models that see earlier rounds can pick a move by reasoning about the history, which would make play less random and could resemble coordination. We want the models’ intrinsic behavior.
- *Retry on unparseable output*. If a reply cannot be parsed as R, P or S, the model gets up to two more tries; if all fail, the move is recorded as `None`.

3.3 Detector and Shuffle-Based Threshold

For the Contribution 2 and Contribution 3 we use the Bonjour et al. [1] detector directly: it computes γ and net influence Γ for each pair from the logged moves [(“R”, “S”, “P”), (“R”, “R”, “R”).....] and flags a pair when net influence on each other reaches $\alpha = 0.05$. For the threshold (Contribution 1), we first run this detector on the real moves from say config A, then build a new threshold from shuffled data and a new score, described below.

Step 1: break the coordination. We shuffle each agent’s list of moves independently. This keeps how often each agent plays R, P or S but destroys which moves were played together in the same round. (If a state were used, we would shuffle only within each state, because the detector

compares moves within a state, and mixing moves across states would change those frequencies and could create or hide dependence.) The detector is then run on the shuffled moves, giving net influence in all six directions ($i \rightarrow j$ and $j \rightarrow i$ for each of the three pairs).

Step 2: turn each run into one score.

For each configuration, the detector gives net influence in six directions, two for each of the three agent pairs. The pairs are named by player position: for example, in the configuration (Gemma, Mistral, Qwen), pair (i, j) is Gemma and Mistral, pair (i, k) is Gemma and Qwen, and pair (j, k) is Mistral and Qwen. For each pair, we take the smaller of its two directional values, $\text{net}(i \rightarrow j)$ and $\text{net}(j \rightarrow i)$, because a pair is flagged as colluding only when both directions exceed the threshold, so the weaker direction decides. We then take the largest of these three pair values. This picks the pair that looks most like colluding, so the score answers: how collusive does the most suspicious pair in this configuration look?

Step 3: set the threshold. We repeat Steps 1 and 2 1000 times per configuration, giving 1000 scores. Since shuffled runs contain no coordination, these scores show how high the score gets by chance alone. Now we have 1000 new scores for each configuration. The maximum of these scores represents the absolute ceiling up to which net influence can rise purely from random moves, which could serve as the threshold. However, we do not take the maximum. We take the 99th percentile instead, because the maximum depends on one extreme shuffle and varies a lot between repeats; the percentile ignores such outliers.

3.4 Experimental conditions

We run three conditions. They differ only in the text added to the shared base prompt (prompt structure in the Appendix).

1. No-collusion instructions. Each player gets only the base prompt asking it to choose randomly. We run two versions, one with no history and one where the previous round’s moves of all three players are added, and we combine them in the results. This tests whether the detector falsely flags LLM agents that were never told to collude (Contribution 1).
2. History-mapping collusion. Each player sees the previous round’s moves. Player 2 is told to either apply a fixed mapping to its own last move or copy Player 3’s last move. Player 3 is told to help Player 2 and to guess which of the two Player 2 will probably use, then play the move that gives Player 2 a point. This tests collusion that is coordinated through a shared rule over the history. Because the prompts make both players’ moves a fixed function of the previous round, this breaks the detector’s net influence equations (Contribution 2).
3. Strict collusion (positive control). There is no history. Before each round, a character $c \in \{R, P, S\}$ is drawn uniformly at random. Player 2 is told to play exactly c . Player 3 is told to help Player

2 win, ignoring its own score, and to play the move that c defeats, so Player 2 gets a point whatever Player 1 does. Player 1 gets only the base prompt. We treat the intended collusion probability as 1, because Player 3 is explicitly told to play the move that gives Player 2 a point. With the shuffle-based score and threshold, the instructed pair (Player 2 and Player 3) is flagged in every configuration, since its score (1.03, 0.10 and 0.39) is far above each configuration’s threshold (about 0.005), while the other pairs stay below it.

4. Results

4.1 No-collusion instructions. No pair was flagged in any configuration, with or without the previous round shown as state. Table 1 gives the most dependent pair’s net dependence score for each config.

Config	No State	With previous round as state
A: Gemma, Gemma, Gemma	0.0003	0.0006
B: Gemma, Gemma, Mistral	0.0017	0.0057
C: Gemma, Mistral, Qwen	0.0008	0.0245

Table 1. Collusion score for LLMs given no collusion instruction. For each pair of players we take the smaller of its two net influence values, and the collusion score is the largest of these values across the three pairs, so it measures the pair that looks most like colluding. A pair is flagged when its score reaches $\alpha = 0.05$. No score does.

Table 1 shows that in both settings, with no history and with the previous round shown, and with no collusion instruction given, the Bonjour detector correctly concluded that no pair was colluding. In all six runs, the highest pair net influence stays below $\alpha = 0.05$, so no pair is flagged.

4.3 Strict collusion

Config	Instructed pair (P2–P3)	Net influence	Shuffle null p99
--------	-------------------------	---------------	------------------

A: Gemma, Gemma, Gemma	Gemma–Gemma	1.0328	0.0051
B: Gemma, Gemma, Mistral	Gemma–Mistral	0.1047	0.0056
C: Gemma, Mistral, Qwen	Mistral–Qwen	0.3916	0.0051

Table 2. Strict collusion. Net influence of the instructed pair (the smaller of its two directions) against the 99th percentile of the net influence scores from 1000 shuffled runs. Each configuration’s instructed pair is the only one flagged, by $\alpha = 0.05$ and by the shuffle threshold.

In every configuration, all net influence values for the non-instructed pairs (P1–P2 and P1–P3), in both directions, are below the calculated shuffle threshold and below $\alpha = 0.05$. Their largest single value is 0.00128, against thresholds of 0.0051 to 0.0056.

5. Discussion and Limitations

Our results show that the Bonjour et al. detector can be used on LLM agents, and that a threshold built from each configuration’s own shuffled moves (the 99th percentile of 1000 shuffles) gives a clear detection limit for LLM agents. This threshold is easy to compute from the logged moves alone. It needs no access to model internals or messages, and it can be recomputed for any new set of agents. In our experiments, the detector raised no false alarm when LLMs were not told to collude, and with the shuffle threshold it found the colluded pair in all three configurations.

In a multi-agent system where three LLMs work on a task independently, collusion between two of them can quietly degrade the efficiency of the whole system. Our approach combined with net-influence detection of Bonjour et al. [\[1\]](#) can potentially find the colluding pair in the system.

Limitations

Collusion was prompted, not emergent. We introduced collusion explicitly through prompts. In real settings, collusive behavior can arise naturally from how models respond to each other, and we did not address that threat model. Our results show that the detector can find collusion we instructed. They do not show that it would find collusion that emerges on its own, which may be weaker, or spread across more agents (here only behaviour of two collusion pairs studied). If

emergent collusion looks different from instructed collusion, the thresholds and scores we report may not transfer.

Shared architecture and reasoning were not studied. Agents can behave alike because they share an architecture and training, not because they collude. We did not study this separately. Our three-Gemma configuration is only a partial check, and Gemma’s strong preference for Scissors shows that shared habits can be large. We also turned reasoning off, so we did not test models that reason. Similar reasoning models given the same situation can reach the same choice and look like collusion, and our method has no way to separate that from real collusion. Our no-false-alarm result applies only to the models, prompts and settings we tested.

Future Work

The next step is to run the shuffle-based threshold across more model configurations and check whether a similar pattern holds, with the aim of finding a threshold that works across many configurations instead of one per configuration. We would also let the models reason, and measure to what extent this produces false alarms of collusion, since similar reasoning models can reach the same choice without colluding

6. Conclusion

We tested whether the collusion detector of Bonjour et al., which was built for scripted agents, also works on LLM agents playing three-player Rock-Paper-Scissors. When the models were not told to collude, the detector found no colluding pair, with or without the previous round shown. When we told Player 3 to help Player 2 win, the detector flagged exactly that pair in all three model configurations, but the strength of the signal differed a lot between them (1.03, 0.10 and 0.39). It also has a blind spot. When both players’ moves are fully fixed by the previous round, their moves stop looking dependent to the detector, and the instructed pair went unflagged.

Because LLMs play probabilistically and share habits (Gemma plays Scissors most of the time), a constant threshold may not suit every set of agents. We therefore set the threshold from each configuration’s own data: we shuffled the agents’ moves 1000 times to break any coordination and took the 99th percentile of the resulting scores, which was about 0.005 in all three configurations. The instructed pairs scored far above it and the other pairs scored below it. For AI safety, this suggests that a collusion check on a multi-agent system should report its score next to a noise floor measured on that system’s own logs. Our evidence comes from one run per configuration, in a toy game, with collusion introduced by prompts, so the next steps are repeated runs, collusion that emerges on its own, and groups of more than two colluders.

Code and Data

- **Code repository:** <https://github.com/GovindhKishore/collusion-detector-rps>

References

1. Trevor Bonjour, Vaneet Aggarwal, and Bharat Bhargava (2022). Information-Theoretic Approach to Detect Collusion in Multi-Agent Games. In *Proceedings of the 38th Conference on Uncertainty in Artificial Intelligence*, PMLR 180:223–232, 2022.
<https://proceedings.mlr.press/v180/bonjour22a.html>

Appendix A: Prompt structure

Every prompt starts with “You are Player i.” followed by the shared base prompt.

- **[BASE]**: play one round of Rock-Paper-Scissors, choose randomly, reply with one letter.
- **[HISTORY]**: “Last round these were the moves: Player 1: x, Player 2: y, Player 3: z.”
- **e**: a character from {R, P, S}, drawn at random each round (strict collusion only).

1. No collusion, no history

- Players 1, 2, 3: [BASE]

2. No collusion, with history

- Players 1, 2, 3: [BASE] + [HISTORY]

3. History-mapping collusion

- **Player 1**: [BASE] + [HISTORY]
- **Player 2**: [BASE] + an instruction to either apply a mapping to its own last move ($R \rightarrow P$, $P \rightarrow R$, $S \rightarrow S$) or copy Player 3’s last move, switching between the two + [HISTORY]
- **Player 3**: [BASE] + an instruction to help Player 2 and guess which rule Player 2 will use this round + [HISTORY]
 - If it expects the mapping, play by Player 2’s last move ($R \rightarrow R$, $P \rightarrow S$, $S \rightarrow P$).
 - If it expects copying, play the move that defeats its own last move ($R \rightarrow S$, $P \rightarrow R$, $S \rightarrow P$).
- **Round 1**: there is no history, so [HISTORY] is left out and Players 2 and 3 are told to choose randomly.

4. Strict collusion (no history)

- **Player 1:** [BASE]
- **Player 2:** [BASE] + an instruction to play c
- **Player 3:** [BASE] + an instruction to help Player 2 by playing the move that c defeats
($R \rightarrow S$, $P \rightarrow R$, $S \rightarrow P$)

LLM Usage Statement

We used Claude to help write code, mainly the plotting and visualization scripts and the code written before the sprint, and to help develop the scoring and shuffle-based threshold. We also used it to correct typos and grammar and to tidy our own draft text of this report. The experimental design, the experiments, and the results are our own, and we checked all reported numbers against the logged outputs.