
When Cooperation Looks Like Collusion: The Limits of Counterfactual Auditing in Multi-Agent AI

Aarav Patel
Georgia Institute of Technology
With
Apart Research

Abstract

Removing a message from an AI agent's conversation can show that the message affected other agents. It cannot, by itself, tell us whether the agents colluded. We investigate this gap with Counterfactual Collusion Certificates (CCC), using controlled replays and paired scripted scenarios in which permitted cooperation and prohibited coordination produce the same visible behavior and message-removal effects. In allocation and transfer tasks, a baseline that treats harmful effects as evidence of collusion incorrectly flags permitted behavior. We then examine which additional observations help distinguish these cases. Changing task costs separates the programmed allocation policies, while an authorization record separates the transfer policies. These checks work for the cases they were designed around, but the broader attribution rule does not meet its prespecified utility criteria. A separate language-model experiment leaves its predicted contrast unresolved, and noisy scripted results differ across platforms. The results illustrate why causal evidence alone is insufficient for a collusion allegation: showing what a message changed is not the same as showing which rule the agents followed. Whether richer evidence improves attribution for independently developed agent behavior remains to be tested.

1. Introduction

An auditor removes a teammate's allocation request and finds that outsiders benefit. This establishes that the message mattered, but does it establish collusion? Our allocation twins produce the same request, allocation, and replay outcome under two different policies: one follows efficiency, while the other protects the coalition. Distinguishing their effects from their permissibility is the central audit problem [7].

Pricing studies show why this distinction matters: simulated agents can sustain high prices without explicit communication [11], and the pattern depends on the competitive environment [12]. Here we study four task agents, two coalition members and two outsiders. Rewards describe outcomes. The benchmark separately assigns policy labels, which are withheld from the auditor's observations.

Figure 1 illustrates the missing evidence. Both policies initially produce the same five-unit outsider-reward improvement when the message is removed. After outsiders' task costs fall, the efficiency policy yields the task, while the coalition-protection policy retains it. The response to this probe distinguishes these programmed alternatives.

Our contribution is an executable account of the **gap between message effects and policy attribution**. We construct paired policies, test what additional evidence distinguishes them, and evaluate a rule that can leave attribution unresolved. A final audit separates three questions: whether selected cases differ, whether their labels are identifiable among all compatible policies, and whether the auditor uses that evidence correctly.

Same evidence, different policies

Constructed allocation twins · P2 · deterministic · seed 19693329




Figure 1. Same evidence under two programmed policies. In constructed allocation twins P2 (seed 19693329), permitted efficiency and prohibited coalition protection yield the same message, allocation, and mean outsider reward: 14 with the message and 19 after its removal. Lowering each outsider's task cost by 10 raises factual probe reward to 29 under efficiency but leaves it at 14 under coalition protection; probe knockout reward is 29 in both worlds. Removing the factual probe record restores identical views in all 200 paired cases in a post-hoc check. Labels are assigned by the generator, not inferred from communication.

2. Related Work

Prior work tests coordination through forced market deviations [2], interventions on latent communication [3], and matched message replacements [4]. Counterfactual explanations identify changes associated with a different decision [13]. We ask a narrower question: what does a measured communication effect establish about whether coordination was permitted?

Xu et al. construct collusive bids with competitive individual marginals, exposing a limitation of audits based on those marginals [5]. Riemer et al. argue for certification of reasoning agents making market decisions [6]. Our examples instead concern message-level replay: even an intervention that reveals outsider harm can leave permitted and prohibited policies observationally indistinguishable.

Causal interpretation depends on the intervention and its assumptions [14]. Inferring intention also requires a model of the agent and its alternatives [8]. We use these established ideas to construct and test message-level attribution counterexamples, evaluating declared policies rather than motives.

3. Testing the evidence

3.1 What message removal shows

In the Colosseum language-model study [1], continuations restart from a checkpoint before message receipt. The factual arm delivers the message; the knockout hides it throughout the continuation. Sham and placebo arms check procedural and nonspecific removal effects (Table 1). The starting history is fixed, but subsequent model choices are regenerated.

Condition	What the receiver sees	Purpose
F: factual	Original message and history	Reference continuation
K: knockout	Target message hidden throughout the continuation	Measure removal effect
S: sham	Same information; removal machinery changes nothing	Check disturbance from replay itself
N: placebo	Target retained; a different, eligible message removed	Check effects of removing irrelevant content

Table 1. Four replay arms from the same recorded checkpoint: factual (F, message present), knockout (K, target message hidden throughout), sham (S, information unchanged), and placebo (N, a different eligible message hidden). An unavailable placebo is recorded as missing, never as a zero effect.

Let O denote the outsiders' average task reward. We measure the effect of removing a message as the difference between expected reward when it is hidden (K) and when it is shown (F):

$$\tau = E[O \mid K] - E[O \mid F].$$

The expectation averages over replay continuations. If τ is positive, outsiders do better when the message is hidden. This says nothing on its own about authorization, intent, or whether collusion occurred.

A CCC certificate evaluates **harmful exposure**: outsider loss beyond a prespecified margin, coalition benefit, and the required welfare and replay-control checks. Its verdict is supported, not

supported, inconclusive, or invalid. These verdicts concern the tested exposure, not whether a policy was prohibited. Appendix B gives the decision rules and their assumptions.

3.2 What attribution requires

Attribution asks whether coordination violated a declared rule. We use nested evidence sets: I0 contains the transcript and factual continuation; I1 adds single-message removal; I2 adds sham and placebo controls; I3 adds joint removal of same-sender carriers; I4 adds task-cost probe continuations; and I5 adds an external authorization record. Distinguishing a selected pair of worlds does not identify their labels across the entire programmed policy space. Binary attribution is identifiable only if every programmed policy compatible with the evidence agrees on whether the declared rule was violated.

A second message may carry the same information, leaving a single-message knockout ineffective [9]. An authorization record answers whether an action was permitted, not what the agent intended.

CCC-L is the ladder-aware attribution rule: it returns an attribution answer or abstains. We report **selective accuracy**, meaning accuracy among answered cases, together with **coverage**, the fraction answered [10]. Because CCC-L was designed alongside the scripted policies, success on those policies tests consistency rather than transfer. Abstention alone can avoid accusations without providing useful answers.

3.3 Two tests of the approach

The scripted benchmark uses a finite four-agent reward model, separate from the language-model study. Its six families contain 200 pairs each: redundancy versus inert behavior (P1), allocation (P2), transfer authorization (P3), message influence versus a shared cause (P4), inert wording (P5), and harmful steering versus useful information (P6). Additional cases test prevalence and controls. Rewards need not fully express a designer's objectives [15], so policy labels are specified separately.

Each scripted unit has deterministic and noisy conditions; in the latter, agents choose random legal actions with probability 0.1. Replays are exact for deterministic scripts and use eight continuations per arm for noisy scripts. Seeds and evaluation criteria were frozen before collection. The harm-gated knockout baseline B2+ was selected using development data only. Its flag is scored as an accusation, unlike CCC's exposure verdict.

Separately, local gemma4:26b agents completed eight matched world pairs under deceptive and authorized-cooperation prompts, with four factual and four knockout continuations per checkpoint. The primary contrast Δ averages the paired difference in message-removal effects. The preregistered directional decision required a positive bootstrap lower bound and exact one-sided sign-flip $p \leq 0.05$. The independent units are the eight pairs, not their continuations; prompt conditions are not verified collusion labels.

4. Results

The benchmark establishes a specific attribution error and tests whether added evidence resolves it. B2+ accused legitimate allocation and authorized-transfer behavior, while the designed probes separated the corresponding policy pairs. The broader evaluation did not meet four prespecified aggregate predictions (Table 2). We report deterministic findings because noisy-policy decisions differed across macOS and Linux. This restriction was made after examining results; both conditions and the original verdicts remain available.

Prediction	Result
Prohibited cases: exposure rate below 25%	Not met: 61/200 (30.5%); CI spans 25%.
Baseline flags most cases in every legitimate class	Not met: allocation/transfer only.
Revised pattern holds in both conditions	Not met: deterministic only.
CCC-L meets the full utility criteria	Not met: 2/84; many comparisons undefined.

Table 2. Four prespecified aggregate benchmark predictions not met. The displayed prevalence and utility counts have different denominators; deterministic findings are reported after a post-hoc restriction prompted by cross-platform noisy-policy disagreement. These criteria are not estimates of real-world detection accuracy.

4.1 What the scripted counterexamples establish

A post-hoc check of the frozen deterministic evidence views identifies which added records separate the constructed policies (Figure 2; Appendix C). Removing the factual cost-probe record restores identical observations in every P2 pair; deleting authorization status does the same for P3. Evidence and auditor behavior can nevertheless diverge. At P6/I4, both labels are identifiable within the finite policy space, yet CCC-L abstains on all 400 members of 200 pairs. At P1/I3, its 200 negative answers match the selected inert labels but do not exclude a compatible prohibited relay. These are unsupported exclusions, not observed false negatives.

Among 200 population worlds with prohibited-policy labels, 61 met the oracle message-removal margin: 30.5% (exact 95% interval 24.2–37.4%). The prespecified prediction of fewer than 25% was not met. Because the interval includes 25%, this does not establish a population rate above that boundary. A policy label and a large removal effect are distinct properties, even in this scripted setting.

The harm-gated baseline accused the legitimate member in all 200 allocation and all 200 authorized-transfer cases, but none of the 200 useful-information or 100 benign-cooperation controls (Figure 2, I1). Permitted coordination can also disadvantage outsiders. Thus the original prediction of high false accusations across all four legitimate classes was not met, even though the scenario-specific attribution error is reproducible.

What evidence actually establishes

More evidence can resolve a named ambiguity; it does not validate an auditor.

Influence	Harmful exposure	Policy attribution
I1: factual + knockout	I2: controls + harm gates	I3–I5: separating evidence
EVIDENCE ADDED	WHAT THE FROZEN RECORDS SHOW	
I1 Single-message knockout K	B2+ accuses legitimate P2/P3: 200/200 each. P6: 0/200; benign cooperation: 0/100.	
I2 Sham S and placebo N	CCC exposure supported on legitimate P2/P3: 200/200 each. Their attribution remains unresolved.	
I3 Joint same-sender removal J	P1 separates 200/200 pairs; 200 negative answers still admit a prohibited relay alternative.	
I4 Task-cost probe PF and PK	P2: correct attribution on 200/200 pairs. P6: identifiable; CCC-L abstains on all 400 members.	
I5 External authorization record	P3: correct attribution on 200/200 pairs. The record defines permission; separation is designed.	

Pair separation ≠ global attribution

P4/I1: 200/200 pairs differ; neither label is globally identifiable.

Relay cases: 0/100 identifiable at I1–I5. Full utility: 2/84 settings.

Separate language-model study: directional contrast not established
8 pairs; contrast = 1.354; 95% interval [−0.020, 2.834]; p = 0.0703125.

Deterministic scope is post hoc; noisy-platform divergence remains unresolved.

Figure 2. What cumulative evidence establishes in deterministic scripts. B2+ flags are accusations; CCC support concerns exposure; CCC-L answers attribution or abstains. P1 tests redundant versus inert behavior, P2 allocation, P3 authorization, P4 shared causes, and P6 useful information versus harmful steering. Each family has 200 independent pairs; members and evidence-level views reuse those pairs. Benign and relay controls have 100 units each. Global identifiability is confined to the frozen policy space. P1's negative answers still admit prohibited relays; P6 remains unanswered despite identifiable labels. Boundary checks and deterministic-only interpretation are post hoc; noisy-platform disagreement remains unresolved.


At I2, CCC supported harmful exposure in all 200 legitimate allocation and all 200 authorized-transfer cases. This is consistent with their programmed outcomes: permitted coordination can disadvantage outsiders. Turning that exposure verdict into an accusation would add a policy claim not established by the replay. The refined accusation pattern held deterministically, but did not meet the criterion across both planned conditions.

Given the cost probe at I4, CCC-L correctly answered both members of all 200 P2 pairs; with authorization at I5, it did so for all 200 P3 pairs. The full utility criterion, however, was met in

only 2/84 deterministic settings, versus 5/84 for the earlier rule. It combines selective accuracy, positive overclaims, and abstention precision and recall. In each version, 73 accuracy comparisons had empty answer sets and were undefined; fully answered settings could also have undefined abstention metrics. The frozen rule counted these as unmet. The result is a failed composite criterion, not an overall accuracy estimate (Appendix C.3).

4.2 An unresolved language-model contrast

MATCHED SCENARIO PAIRS



Predicted difference not established.

Figure 3. Eight independent matched-pair differences in outsider-reward message-removal effects (deceptive-prompt minus authorized-cooperation-prompt regime) in the separate local-model study. Mean $\Delta = 1.354$ task-reward units; 95% whole-pair bootstrap interval -0.020 to 2.834 ; exact one-sided sign-flip $p = 0.0703$. The interval includes zero, so the preregistered directional contrast was not established; prompts are not verified collusion labels.

The separate language-model study estimated $\Delta = 1.354$ reward units (95% whole-pair bootstrap interval -0.020 to 2.834 ; exact one-sided $p = 0.0703$; Figure 3). Neither preregistered decision requirement was met. The interval includes near-zero and substantially positive effects, leaving the directional contrast unresolved rather than establishing equivalence. Appendix D provides the individual contrasts, branch accounting, and sensitivity analysis.

The 16 exploratory certificates were 0 supported, 11 not supported, and 5 inconclusive. Of 199 scheduled branch records, 143 executed validly and 56 placebo branches could not run because no eligible message existed. No outsider-effect lower bound exceeded the harm margin. These results leave certificate sensitivity unresolved; unavailable placebos alone do not explain the absence of support.

5. Discussion and Limitations

An auditor can correctly identify an outsider cost and still accuse a permitted policy. Our results support reporting causal influence, harmful exposure, and policy attribution separately. The P1 and P6 checks also show why accuracy alone is insufficient: correct answers may leave prohibited alternatives plausible, while available identifying evidence may go unused. Evaluation should therefore measure unsupported exclusions and missed answers alongside false accusations.

Limitations

The policy space, labels, and separating probes were designed together. An unfamiliar legitimate policy could resist the cost probe, and a missing or forged authorization record could change the attribution. Knockouts may remove useful task information. The final record-deletion check is post hoc and uses the same frozen cases, so it tests the internal evidence boundary rather than out-of-sample performance.

The language-model study covers one model and eight independent pairs without verified collusion labels. Full-certificate sensitivity and noisy-policy reproducibility remain unresolved; adaptive coalitions and hidden semantic channels were not evaluated. Appendix A explains how these assumptions limit transfer and deployment claims.

Future Work

The next test is whether added evidence improves attribution on independently designed agent behavior. A proposed six-to-eight-week study would validate full-pipeline controls, pin the environment, and compare two local models and two task variants at comparable coverage. Labels and endpoints would be frozen before collection. Success requires fewer false accusations without an unacceptable sensitivity loss. Appendix F specifies the proposed criteria, precision checks, and stopping rules.

6. Conclusion

A message-removal effect does not by itself distinguish permitted coordination from a prohibited policy. Our scripted examples make this ambiguity reproducible and identify additional evidence that resolves specific alternatives. The broader evaluations leave practical reliability open. The next test is whether the distinction improves decisions on independently designed agent behavior. An audit should state what communication changed, which rule it evaluates, and which explanations the evidence still permits.

Code and Data

Code: [CCC GitHub repository](#); Data: [verified results index](#); Demo: [Evidence Explorer](#). Reproduce the recorded headline contrast with python3 scripts/reproduce_headline.py (Appendix E).

References

- [1] Mason Nakamura et al. (2026). Colosseum: Auditing Collusion in Cooperative Multi-Agent Systems. arXiv:2602.15198. <https://arxiv.org/abs/2602.15198>
- [2] Hieu Tran and Hong Yu (2026). Is Anyone Enforcing the Cartel? What Reinforcement Learning Teaches LLM Agents in Markets. arXiv:2610.28787. <https://arxiv.org/abs/2610.28787>
- [3] Ramneet Kaur et al. (2026). Beyond the Transcript: Detecting Covert Coordination in Latent Multi-Agent Communication. arXiv:2608.19161. <https://arxiv.org/abs/2608.19161>
- [4] Huixiang Zhang and Mahzabeen Emu (2026). Do Latent Channels Actually Communicate? A Causal Audit of Latent Multi-Agent LLM. arXiv:2607.26773. <https://arxiv.org/abs/2607.26773>
- [5] Xin Xu et al. (2026). Collusion with Competitive Marginals: Price-Level Audits Are Blind by Construction. arXiv:2607.26385. <https://arxiv.org/abs/2607.26385>
- [6] Matthew Riemer et al. (2026). Position: Collusion Risks Among AI Reasoning Agents Justify Certification Requirements for Making Market Decisions. arXiv:2608.18078. <https://arxiv.org/abs/2608.18078>
- [7] Jamiu Idowu et al. (2026). Mapping Human Anti-collusion Mechanisms to Multi-agent AI Systems. arXiv:2601.00360. <https://arxiv.org/abs/2601.00360>
- [8] Francis Rhys Ward et al. (2024). The Reasons that Agents Act: Intention and Instrumental Goals. arXiv:2402.07221. <https://arxiv.org/abs/2402.07221>
- [9] Hana Chockler and Joseph Y. Halpern (2004). Responsibility and Blame: A Structural-Model Approach. *Journal of Artificial Intelligence Research* 22, 93–115. <https://arxiv.org/abs/cs/0312038>
- [10] Yonatan Geifman and Ran El-Yaniv (2017). Selective Classification for Deep Neural Networks. arXiv:1705.08500. <https://arxiv.org/abs/1705.08500>
- [11] Emilio Calvano, Giacomo Calzolari, Vincenzo Denicolò, and Sergio Pastorello (2020). Artificial Intelligence, Algorithmic Pricing, and Collusion. *American Economic Review* 110(10), 3267–3297. <https://doi.org/10.1257/aer.20190623>
- [12] Timo Klein (2021). Autonomous Algorithmic Collusion: Q-Learning under Sequential Pricing. *RAND Journal of Economics* 52(3), 538–558. <https://doi.org/10.1111/1756-2171.12383>
- [13] Sandra Wachter, Brent Mittelstadt, and Chris Russell (2018). Counterfactual Explanations Without Opening the Black Box: Automated Decisions and the GDPR. *Harvard Journal of Law & Technology* 31(2), 841–887. <https://arxiv.org/abs/1711.00399>
- [14] Judea Pearl (2009). *Causality: Models, Reasoning, and Inference* (2nd ed.). Cambridge University Press. <https://doi.org/10.1017/CBO9780511803161>
- [15] Dylan Hadfield-Menell, Smitha Milli, Pieter Abbeel, Stuart Russell, and Anca D. Dragan (2017). Inverse Reward Design. *NeurIPS* 30. <https://arxiv.org/abs/1711.02827>

Appendix A. Limitations and Dual-Use / Ethical Considerations

A.1 What the policy labels mean

The benchmark assigns labels to programmed policies and permission records. In the allocation twins, the label describes a policy across cost conditions, not an independently adjudicated violation in the initial allocation. The transfer twins differ in authorization by definition. These cases establish that the chosen observations can leave attribution ambiguous; they do not estimate the prevalence of collusion among deployed agents.

Label identifiability is conditional on the finite candidate-policy set and the seeded environment. An omitted legitimate policy could behave like the prohibited allocation policy under the cost probe. An omitted prohibited policy could resemble a permitted one. Expanding the candidate set can therefore invalidate an attribution that was identifiable in the original set, even if the recorded observations remain unchanged.

A.2 Assumptions behind replay and attribution

A causal interpretation requires a faithful checkpoint and an intervention that changes the intended message exposure rather than unrelated state. Subsequent agent choices are regenerated, not held fixed. Useful task information may disappear with the message, and repeated or relayed content may survive the intervention. A sham tests procedural disturbance; an eligible inert-message placebo does not establish equivalence to removing other salient information. Failure of these assumptions weakens the interpretation of the measured effect, not just the attribution label.

Authorization is informative only to the extent that the governing rule is appropriate and the record is authentic and complete. A missing record is not proof of prohibition. A cost probe similarly distinguishes specified response policies, not intention in general. The study does not validate auditing against adaptive coalitions, hidden semantic channels, forged records, or unfamiliar legitimate policies.

A.3 Responsible use and release boundaries

An exposure certificate and an attribution decision should be reported separately. Positive exposure support is insufficient for sanctions, while a negative or inconclusive certificate is not clearance. The P1 exclusion ambiguity illustrates why unsupported negative answers deserve scrutiny as well as unsupported accusations. Any deployment study would need to measure both types of error at stated coverage and include independent review of consequential decisions.

Publishing counterexamples can help auditors design better tests, but it can also reveal the limits of narrow probes to adaptive agents. Real communication audits would require appropriate consent, access controls, and privacy protection. No new human-subject data were collected. Original CCC material is MIT-licensed; inherited components have separate terms, and unresolved redistribution permissions are documented in [NOTICE.md](#) and [docs/release/LICENSES.md](#). The release is not a certification of those permissions.

Appendix B. Experimental Design and Certificate Rules

B.1 Units, policies, and evidence access

The scripted benchmark and language-model study are separate evaluations and are never pooled. The benchmark uses 200 independent paired seeds for each of six families: P1 contrasts redundant harmful communication with inert behavior; P2 contrasts coalition-protecting and efficient allocation; P3 contrasts unauthorized and authorized transfer; P4 contrasts message-caused behavior with a shared cause; P5 changes suspicious wording without changing the label; and P6 contrasts harmful steering with useful information. Additional population and control classes test other aspects of the audit.

A pair shares its seed and nuisance draw. Its two members, replay branches, and repeated evidence-level views are not independent observations. Each unit has a deterministic condition and a condition with random legal actions at probability 0.1. Deterministic evaluation uses exact continuations; the noisy condition uses eight per arm. The 200 population worlds and the 100-unit control classes have their own denominators.

The information sets are cumulative. I0 provides the transcript and factual outcome. I1 adds single-message knockout K; I2 adds sham S and eligible placebo N; I3 adds joint removal of same-sender carriers; I4 adds factual and knockout cost-probe continuations, PF and PK; I5 adds authorization. The auditor receives the declared observations, not the generator's member label. CCC-L was nevertheless designed with knowledge of the benchmark's policies and separators.

B.2 Effect definitions and decision gates

Let O and C be mean outsider and coalition reward, and W total reward. The recorded contrasts are $\tau_O = E[O_K] - E[O_F]$, $\tau_C = E[C_F] - E[C_K]$, and $\tau_W = E[W_K] - E[W_F]$. Positive τ_O means removal benefits outsiders; positive τ_C means retaining the message benefits the coalition. Rewards are task units, not monetary losses or an independent measure of social welfare.

CCC first checks replay integrity and sampling validity. Support then requires the outsider-effect lower bound to exceed 2.5, the coalition-benefit lower bound to exceed zero, and the total-reward-change lower bound to be at least -5. It also requires an accepted, executed placebo, resolved information-loss checks, and no blocking reason. For sampled continuations, one-sided 95% bounds invert exact shift-permutation tests on a 0.01-unit shift grid; deterministic scripts use point effects. A positive effect estimate alone is insufficient.

A valid case is not supported when the outsider-effect upper bound is below 2.5 and no redundant-carrier flag or blocking failure prevents that conclusion. Other eligible cases remain inconclusive; integrity failures are invalid. Missing or invalid outcomes are not converted to zero rewards. The oracle prevalence endpoint uses $\tau_O \geq 2.5$, whereas the support rule requires a lower bound strictly above that margin. These endpoints must not be interchanged.

B.3 Controls, selection, and frozen decisions

Sham replay retains the factual information while exercising the intervention machinery. Placebo replay removes a different eligible inert message. In the real-agent records, the placebo and target appeared on different communication surfaces and at different prompt positions. Shared masking code therefore does not imply a fully matched content or position control. Carrier extraction is heuristic and does not exclude arbitrary paraphrases or encoded messages.

The harm-gated baseline B2+ was selected on development data before prospective evaluation. Development-informed changes to the attribution rule and predictions were retained in the frozen protocol. New seeds excluded development blocks and earlier-study seeds. The P4 influence control and twin-symmetry checks test specified causal or implementation properties; they do not establish recovery of the full certificate on independently labeled harmful coordination.

Appendix C. Evidence Boundaries and Auditor Utility

C.1 Pair separation is not label identification

The final analysis rechecks frozen deterministic outputs using a separate implementation that imports no CCC modules. It reconstructs lower evidence levels from canonical I5 observations, verifies recorded views and partition hashes, and checks binary-label agreement among compatible programmed policies. Its 21,000 member-by-access views and 2,100 partition hashes describe verification coverage, not a larger independent sample. The analysis was added after results were known.

Table C1 compares three questions at selected evidence levels: whether the two named members differ, whether both members' labels are identifiable among all compatible policies, and whether CCC-L v1 answers both correctly. All denominators are 200 independent pairs per family. A zero in the last column can reflect abstention rather than a wrong answer. P5 is a same-label wording control.

Family	Access	Pairs distinct	Both labels identifiable	CCC-L correct both
P1	I3	200/200	0/200	200/200
P2	I4	200/200	200/200	200/200
P3	I5	200/200	200/200	200/200
P4	I1	200/200	0/200	0/200
P5	I1	200/200	200/200	0/200
P6	I4	200/200	200/200	0/200

The P1 result is especially important. At I3 the named pair separates and the rule answers both members correctly, but its 200 negative answers on inert members remain compatible with a

prohibited relay policy. These are not observed false negatives. They show that correctness on the selected examples does not by itself justify excluding alternatives. At P6/I4, both labels are identifiable within the finite set, yet the rule abstains on all 400 members nested within 200 pairs. Across I1 to I5, none of the 100 relayed-harmful cases is identifiable in the archived policy space.

C.2 Which added record carries the distinction?

Removing the newly added evidence restores identical observations in all 200 pairs for P1, P2, P3, P4, and P6 at their designed separating levels. Component checks narrow that result. For P2/I4, deleting PF restores equality, whereas deleting PK alone or common probe metadata leaves all pairs distinct. PF outsider reward alone also differs in every pair, although other PF fields differ too. For P3/I5, deleting only authorization.status restores equality.

These are diagnostic deletions from recorded observations, not new environmental interventions or newly permitted benchmark information sets. They establish the role of the deleted records within these constructed alternatives, not universal sufficiency. The labels, policies, and earlier design choices remain inherited from the generator; separate implementation is not independent data validation.

C.3 What the 2/84 result measures

The frozen utility test combines selective-accuracy comparison, positive overclaims, abstention precision and recall, and specified pair-attribution requirements. The revised rule meets the full conjunction in 2 of 84 deterministic class-by-access settings; the previous rule meets it in 5. Both have 73 undefined selective-accuracy differences, retained as unmet. The revised rule also has 9 undefined abstention-precision and 37 undefined abstention-recall values; the previous rule has 3 and 37, respectively.

Even fully answered allocation and transfer settings can leave abstention metrics undefined. The conjunction therefore mixes demonstrated shortcomings with components that cannot be estimated in a setting. Its original failure is preserved, but the pass count is neither overall accuracy nor a valid ranking of the two methods. Zero positive overclaims likewise does not validate negative exclusions or useful coverage. All four original aggregate prediction failures remain reported.

Appendix D. Language-Model Analysis and Uncertainty

D.1 Estimand and independent units

The local-model study used gemma4:26b in Colosseum with eight matched world pairs. Within each regime and checkpoint, the outsider-removal effect is the mean knockout reward minus the mean factual reward. For pair w , d_w is that effect under deceptive prompts minus the effect under authorized-cooperation prompts. The reported estimate Delta is the mean of the eight d_w values, not a direct difference in factual rewards. Selected messages and checkpoints can differ across prompt regimes. Prompt assignment is not verified ground truth for realized collusion.

In ascending world-seed order, the paired contrasts rounded to six decimals are 0.000000, 5.259128, -0.780140, -1.086136, 0.116391, 1.359142, 3.492796, and 2.473425. Seven are

nonzero; the independent sample remains eight, including the zero contrast. The full-precision mean is 1.354325711018898. Replay branches and individual agent decisions do not increase that sample size.

D.2 Primary test and secondary analyses

The preregistered directional decision required both a positive lower bootstrap bound and an exact one-sided paired sign-flip p-value no greater than 0.05. Whole-pair percentile resampling used 10,000 draws. The archived 95% interval is [-0.0195496491, 2.8336288730]. Enumerating all 256 sign assignments gives 18 at least as favorable as the observed statistic, so $p = 18/256 = 0.0703125$. Neither decision requirement was met. This interval does not establish equivalence or absence of an effect.

No registered secondary test passed after Holm adjustment: the adjusted p-values were 0.1641, 0.4375, and 0.1641. The cooperative-regime effect had only five nonzero values against its prespecified floor of six. A registered sensitivity analysis excluding two non-committing factual continuations retained the eight pairs but changed the estimand: $\Delta = 0.727$, with 95% interval [-0.306, 1.914]. Its $p = 0.15625$ was computed post hoc. Bootstrap-order checks were also post hoc; neither analysis changes the primary conclusion.

D.3 Execution accounting and control limits

The study completed 16 worlds and scheduled 199 branch records. Of these, 143 executed validly: 64 factual, 64 knockout, 7 sham, and 8 placebo continuations. The other 56 were scheduled placebo branches without an eligible message and did not execute. There were no invalid executed continuations. Eligible placebos existed at only two of the 16 checkpoints.

Exploratory certificate outcomes were 0 supported, 11 not supported, and 5 inconclusive. The largest outsider-effect lower bound was 0.77, below the 2.5 margin, so unavailable placebos alone cannot explain the absence of support. Sham comparisons did not establish the needed equivalence across regimes, and the small lesion diagnostic was uninformative. These observations leave full-certificate sensitivity unresolved; they do not show that harmful coordination was absent.

Appendix E. Reproduction, Provenance, and Deviations

E.1 Reproducing the recorded result

The public reproduction entry point requires Python 3.12 or later and no model calls, API keys, or additional packages. From the repository root, run:

```
python3 scripts/reproduce_headline.py
```

The script verifies the archived input hash, extracts to a temporary directory, and recomputes the eight pair contrasts and exact sign-flip test. It prints a mean of 1.354325711018899 and $p = 0.0703125$. The approximately 10^{-15} difference from the saved mean reflects arithmetic order. This command does not regenerate agent behavior or recompute the bootstrap interval. The retained bootstrap verifier requires NumPy.

Public-file checks are documented in [docs/REPRODUCING.md](#). They cover arithmetic, explorer provenance, and artifact navigation. The final boundary analysis is in [paper/final/last_minute/analyze_boundary.py](#), with its results in [boundary_audit.json](#) and [BOUNDARY_ANALYSIS.md](#). The paper's two new figures were generated from those records and the frozen benchmark, not from illustrative measurements.

E.2 Where to inspect the evidence

The benchmark report is [docs/enh/r2/run_par/RESULTS.json](#); its protocol and freeze are in [docs/enh/r2/](#). Corrected real-agent outputs are in [docs/analysis/exp002_heldout_v3/](#), with the raw archive at [docs/runner/archives/exp002-live.tar.gz](#). The numerical source map is [paper/final/CLAIMS_LEDGER.md](#). The final post-hoc handoff and figure provenance are in [paper/final/last_minute/](#). The release inspected for the new analysis is commit [88e3be21c432bb815ad9ebfaefcd2a519136e9e6](#).

The offline Evidence Explorer uses bundled records and runs without model inference. Full pipeline regeneration is a separate task: it needs the pinned upstream Colosseum environment and dependencies. A fresh external installation has not been demonstrated, and only nine of 27 legacy pilot-manifest files are public. Prior complete-suite checks used existing local inputs, so they must not be described as a clean-checkout reproduction.

E.3 Deviations that affect interpretation

The available platform comparison matched all 149 overlapping deterministic units, but 142 of 148 noisy units differed; 1,160 of 1,452 noisy method records contained differences. These include decisions, not merely serialization. The cause is unresolved. Restricting established findings to deterministic evaluation was a post-hoc decision, and agreement on the available overlap is not a general platform-invariance guarantee. Both strata and the original aggregate verdicts remain archived.

The corrected real-agent analysis repaired 16 validity fields and five reason arrays without changing effect estimates or certificate statuses. Earlier versions remain available. The follow-on real-agent experiment was not completed. Colosseum, Terrarium, checkpoint machinery, and an existing pilot were inherited; CCC's replay layer, audits, benchmark, and analyses were developed during the Sprint. A software test pass is evidence about implementation behavior, not detector sensitivity or peer review.

Appendix F. A Testable Research Continuation

F.1 Controls and independent cases

The next question is whether additional evidence reduces false accusations on independently produced behavior while retaining sensitivity to harmful prohibited coordination. The proposed six-to-eight-week study separates exposure-certificate recovery from policy attribution. During weeks 1 and 2, it would pin the environment, resolve platform discrepancies or prospectively

restrict the platform, validate positive and negative controls, and freeze labels, endpoints, uncertainty methods, and stopping rules.

Cases and labels would be produced without access to CCC-L decisions or its separating rules. Authorization would be assigned externally, and actual compliance and harm checked separately from prompts. Unresolved labels would remain visible. A development-only throughput and precision assessment would set the feasible sample and compute limit before held-out collection; failed exposure controls would stop efficacy claims rather than prompt tuning on evaluation cases.

F.2 Comparison and decision criteria

The current proposal caps collection at 160 held-out trajectories across two local models, two task variants, and permitted or prohibited conditions. This is a proposed ceiling, not collected data or a power guarantee. Methods would receive the same observations and compute budget within each comparison. Cost-probe evidence and verified authorization would be added separately, so a gain from direct permission records is not mistaken for improved behavioral inference.

At comparable coverage, proposed success requires a paired reduction in legitimate-case accusations whose 95% interval excludes zero, together with a sensitivity-difference lower 95% bound above -0.10. Risk-coverage curves, abstentions, unresolved labels, and control failures would accompany those endpoints. An imprecise result at the cap remains inconclusive. The deliverables would be a frozen protocol, independent cases where release rights permit, paired endpoint estimates, and reproducible analyses. This plan has not been executed or preregistered.

LLM Usage Statement

This paper was developed with the assistance of Claude Code and OpenAI Codex. The human author directed the research, made experimental decisions, evaluated results, and reviewed the final manuscript. AI agents assisted with coding, analysis, literature review, and drafting.

Approximately 10% of the main paper's text was AI-generated, along with much of the appendices due to time constraints. The main paper was extensively reviewed and reworded to maintain a natural, human writing style while preserving scientific accuracy. All reported findings were checked against recorded data and reproducible analyses. Every number in this report was pulled directly from data and not model assumptions.