

TwinSignal: observational twins and randomized channel probes for agent audits

Jun Kwon | Independent participant

Research conducted during the AI Collusion Research Sprint, October 2026

Abstract

Correlated actions can reflect either communication between agents or a shared environment. Treating statistical dependence as proof of collusion therefore risks accusing agents that never exchange information. We construct an executable synthetic benchmark in which a common-cause mechanism and a directed-copy mechanism produce exactly the same passive action traces. No statistic computed solely from those traces can distinguish the mechanisms. We then randomize the signal delivered through the candidate channel and measure the recipient's response using a one-sided exact Fisher test. A factorial experiment varies sample size, response noise, and intervention uptake across 400 seeds per configuration. With 128 rounds and 10% response noise, passive association is detected in all 400 common-cause and all 400 directed-copy runs. Under full channel intervention, the corresponding rates are 9/400 and 400/400. When intervention uptake is zero, the two mechanisms again coincide. These are controlled binary-policy simulations, not observations of language models or unlawful market conduct. The contribution is a small, reproducible audit fixture with matched negative and positive controls and an explicit failure condition. It demonstrates why an audit should separately report action association, channel responsiveness, and evidence of harmful coordination, rather than treating them as interchangeable findings.

1. Introduction

An auditor sees two agents choose similar actions. This may be suspicious, but it is also what one expects when both receive the same weather forecast, demand estimate, or task instruction. A detector can correctly identify dependence and still support an incorrect explanation of that dependence.

We study an intentionally small question: when can a controlled change to a suspected communication channel distinguish dependence caused by that channel from dependence caused by an unobserved common input? The benchmark is deliberately narrower than detecting collusion. It has no prices, rewards, strategic learning, concealed natural-language messages, or claims about real-world intent.

The contributions are (1) paired generators whose passive outputs are exactly identical although their causal mechanisms differ, (2) a reproducible intervention experiment with exact tests and Monte Carlo uncertainty, and (3) a zero-uptake negative control showing that the intervention cannot identify a mechanism it does not reach. These are an audit fixture and worked example of familiar causal reasoning, not a new identification theorem or a state-of-the-art collusion detector.

2. Related work

Bonjour, Aggarwal, and Bhargava (2022) investigate information-theoretic detection in multi-agent games, including repeated Rock-Paper-Scissors and poker. Their work motivates measuring action dependence. Our synthetic construction addresses a different, narrower issue: a dependence statistic alone does not tell the auditor which causal mechanism generated the trace. We do not reproduce their full net-influence statistic, their games, or their empirical performance. Mutual information here is descriptive; the reported rejection rates use Fisher's exact test.

The current sprint asks what constitutes convincing evidence of collusion and proposes scripted controls as a low-cost starting point. Our experiment makes the distinction between observation and intervention explicit. Comparison against a full game-based method or a language-model monitor remains future work.

3. Methods

3.1 Two observationally identical mechanisms

For every round t , draw an independent shared bit $S_t \sim \text{Bernoulli}(1/2)$ and noise $E_t \sim \text{Bernoulli}(e)$. The first agent's passive action is $A_t = S_t$.

In the common-cause mechanism, $B_t = S_t \text{ XOR } E_t$: both agents respond to the environment and there is no A-to-B channel. In the directed-copy mechanism, $B_t = A_t \text{ XOR } E_t$: the second agent reads the first agent's action through a channel. An independent control instead draws $B_t \sim \text{Bernoulli}(1/2)$.

The two structured mechanisms have the same distribution $P(A,B)$: $P(00)=P(11)=(1-e)/2$ and $P(01)=P(10)=e/2$. More strongly, the paired implementation uses the same S and E arrays, so its passive traces are identical bit for bit. Any deterministic statistic of those traces is consequently identical. A randomized detector also has the same output distribution if its own randomness is independent. This is an observational identification limit, not a failure specific to mutual information.

3.2 Intervention and estimand

On intervention rounds, independently draw a fair bit Z_t and deliver it in place of the first agent's outgoing signal. The common-cause recipient still responds to S . The copying recipient uses Z with probability u and otherwise receives the original S ; the Bernoulli delivery indicators are independent of all other draws. Noise is applied after signal selection.

We test positive association between Z and B . For the copying mechanism,

$$P(B=1 \mid Z=1) - P(B=1 \mid Z=0) = u(1-2e).$$

For the common-cause and independent controls this difference is zero. Full delivery therefore identifies channel responsiveness within this model. It does not identify malicious intent, a welfare loss, or prohibited cooperation. The intervention also assumes an auditor can modify the channel without changing the shared environment or introducing another route to B .

3.3 Experiment and analysis

The preselected grid is n in $\{32,128,512\}$, e in $\{0.1,0.3\}$, and u in $\{1,0.5,0\}$ for intervention conditions. Passive conditions use $u=1$. Each configuration uses seeds 0 through 399. The same seed indexes are deliberately reused across conditions; cells are paired and must not be pooled as independent observations. There are 72 mechanism/configuration cells, each with 400 repetitions.

For binary arrays X,Y , let a count joint ones, r count $X=1$, and c count $Y=1$. Conditional on the margins under independence, the probability of a joint-one count k is $C(r,k) C(n-r,c-k) / C(n,c)$. The one-sided p-value sums these probabilities over $k \geq a$. A positive association is flagged at $p \leq 0.05$. Tests are directional and do not detect arbitrary inverse coordination. We report per-cell rates, not familywise guarantees across the whole grid.

Wilson 95% intervals describe Monte Carlo uncertainty in each rate. They do not quantify model misspecification or uncertainty about real agents. Plug-in mutual information is also reported in bits, without treating finite-sample positive estimates as proof of dependence.

All experiments use Python's standard library and local CPU. The exact-test implementation was cross-checked against `scipy.stats.fisher_exact` on 3,055 contingency tables, with maximum absolute difference $1.1e-16$. A full rerun reproduced the result JSON exactly. This is numerical cross-validation, not independent scientific review. No model inference, paid API, external dataset, or participant data is used. The source includes analytic checks on Fisher probabilities, exact mutual-information examples, deterministic reruns, and equality of the passive twins.

4. Results

Table 1 shows the $n=128$, $e=0.1$ configuration. Counts are positive-association flags out of 400 repetitions. Parentheses contain Wilson 95% intervals in percent.

Audit condition	Independent	Common cause	Directed copy
Passive	13/400; 3.25% (1.91-5.48)	400/400; 100% (99.05-100)	400/400; 100% (99.05-100)
Intervention, $u=1$	13/400; 3.25% (1.91-5.48)	9/400; 2.25% (1.19-4.22)	400/400; 100% (99.05-100)
Intervention, $u=0.5$	13/400; 3.25% (1.91-5.48)	9/400; 2.25% (1.19-4.22)	399/400; 99.75% (98.60-99.96)
Intervention, $u=0$	13/400; 3.25% (1.91-5.48)	9/400; 2.25% (1.19-4.22)	9/400; 2.25% (1.19-4.22)

Passive mean mutual information is 0.53935 bits for each twin. Under full intervention it is 0.00526 bits for the common-cause recipient and 0.53864 bits for the copier. With zero uptake, the copier and common-cause recipient are again identical. The common-cause passive flags are statistically correct dependence detections; calling them false positives is appropriate only if someone incorrectly interprets them as evidence of communication or collusion.

Figure 1 displays the full intervention sensitivity grid; Appendix B reports all 72 cells. At noise $e=0.3$ and uptake $u=0.5$, the association flag rate rises from 74/400 (18.5%) at $n=32$ to 399/400 (99.8%) at $n=512$. This is a power sensitivity within the specified simulator, not a validated performance guarantee for deployed agents. Error bars show per-cell Monte Carlo uncertainty; shared controls are intentionally reused, so the cells are not independent replications of each other. No significance claim for superiority across the grid is made here.




Figure 1. Positive-association rates under randomized signal probes. Lines show 400-seed estimates and vertical bars show Wilson 95% intervals. The zero-uptake copying control overlaps the common-cause control. These are simulated policies, not LLM results.

5. Discussion and limitations

The useful output is a controlled failure case for audits: observational equivalence defeats passive mechanism identification, while a reached and correctly randomized channel can be tested. An audit interface should distinguish “associated actions,” “response to a controlled channel,” and “harmful coordination supported by task evidence.” Our experiment addresses the first two only.

The setup is intentionally simple and its outcome follows analytically from the generator. This limits novelty and external validity. It does not show that any deployed model colludes, that natural-language watermarking is reliable, or that a strategic agent cannot recognize and evade an audit. Rounds are independent; serial correlation, delayed responses, multi-channel communication, adaptive behavior, inversely encoded signals, and

intervention-induced task failure are untested. A non-significant result may mean low power or failed delivery, not absence of coordination. In particular, $u=0$ is an explicit example of a null result with a directed channel in the underlying model.

Future work should validate delivery logs, add benign authorized communication as a positive channel control, test delayed and inverse responses, and measure task utility under interventions. Only then should the fixture be extended to real model trajectories. A stronger collusion claim would also require incentives, harm, and evidence excluding authorized cooperation.

6. Conclusion

The passive twins demonstrate a precise limit: identical logged action distributions cannot reveal whether a common input or a directed channel produced them. Randomized signal probes separate the mechanisms under stated delivery assumptions. The experiment supports careful channel-audit reporting, not an inference of misconduct from correlation.

Code and data

The complete experiment source is included in Appendix C, and every result cell is in Appendix B. There is no required external repository or model download. Source and machine-readable results are also retained locally as `experiment.py` and `results-v001.json`. Reproduce with Python 3.8+ by running `python3 experiment.py --output results.json``; the script refuses to overwrite an existing output.

References

Bonjour, T., Aggarwal, V., and Bhargava, B. (2022). Information theoretic approach to detect collusion in multi-agent games. Proceedings of UAI, PMLR 180:223-232. <https://proceedings.mlr.press/v180/bonjour22a.html>

Apart Research (2026). AI Collusion Research Sprint: research questions, guidelines, and submission template. <https://apartresearch.com/sprints/ai-collusion-research-sprint-2026-10-23-to-2026-10-25>

Appendix: Limitations and Dual-Use / Ethical Considerations

These binary simulations contain no personal data and interact with no market. Channel intervention in a real system would require permission, isolation, and checks that legitimate users are not harmed. A detected response does not establish deceptive or anticompetitive intent. Conversely, publishing detector blind spots can inform evasion; this package demonstrates abstract binary mechanisms and does not provide an operational scheme for concealing real coordination. Test outputs must not be used to penalize people or organizations without independent evidence.

LLM usage statement

Codex proposed the experimental framing, wrote and executed the simulation code, and drafted this report under the participant's instructions. The participant has not been represented as manually writing or independently verifying the manuscript. Reported numerical results come from the local executable experiment. No LLM was used as an experimental agent, and no independent scientific review has yet been completed.

Appendix B. Complete simulation grid

All rates use 400 seeds. Passive conditions have no intervention. CI is the per-cell Wilson 95% interval. Rows share seeds; do not pool them as independent samples.

n	Noise	Probe	Uptake	Mechanism	Flags / 400	Rate	95% CI
32	0.1	no	1.0	independent	7	0.018	0.009-0.036
32	0.1	no	1.0	common_cause	400	1.000	0.990-1.000
32	0.1	no	1.0	directed_copy	400	1.000	0.990-1.000
32	0.1	yes	1.0	independent	18	0.045	0.029-0.070
32	0.1	yes	1.0	common_cause	8	0.020	0.010-0.039
32	0.1	yes	1.0	directed_copy	400	1.000	0.990-1.000
32	0.1	yes	0.5	independent	18	0.045	0.029-0.070
32	0.1	yes	0.5	common_cause	8	0.020	0.010-0.039
32	0.1	yes	0.5	directed_copy	246	0.615	0.566-0.661
32	0.1	yes	0.0	independent	18	0.045	0.029-0.070
32	0.1	yes	0.0	common_cause	8	0.020	0.010-0.039
32	0.1	yes	0.0	directed_copy	8	0.020	0.010-0.039
32	0.3	no	1.0	independent	7	0.018	0.009-0.036
32	0.3	no	1.0	common_cause	264	0.660	0.612-0.705
32	0.3	no	1.0	directed_copy	264	0.660	0.612-0.705
32	0.3	yes	1.0	independent	18	0.045	0.029-0.070
32	0.3	yes	1.0	common_cause	4	0.010	0.004-0.025
32	0.3	yes	1.0	directed_copy	264	0.660	0.612-0.705
32	0.3	yes	0.5	independent	18	0.045	0.029-0.070
32	0.3	yes	0.5	common_cause	4	0.010	0.004-0.025
32	0.3	yes	0.5	directed_copy	74	0.185	0.150-0.226
32	0.3	yes	0.0	independent	18	0.045	0.029-0.070
32	0.3	yes	0.0	common_cause	4	0.010	0.004-0.025
32	0.3	yes	0.0	directed_copy	4	0.010	0.004-0.025
128	0.1	no	1.0	independent	13	0.033	0.019-0.055
128	0.1	no	1.0	common_cause	400	1.000	0.990-1.000
128	0.1	no	1.0	directed_copy	400	1.000	0.990-1.000
128	0.1	yes	1.0	independent	13	0.033	0.019-0.055
128	0.1	yes	1.0	common_cause	9	0.022	0.012-0.042
128	0.1	yes	1.0	directed_copy	400	1.000	0.990-1.000
128	0.1	yes	0.5	independent	13	0.033	0.019-0.055
128	0.1	yes	0.5	common_cause	9	0.022	0.012-0.042
128	0.1	yes	0.5	directed_copy	399	0.998	0.986-1.000
128	0.1	yes	0.0	independent	13	0.033	0.019-0.055

n	Noise	Probe	Uptake	Mechanism	Flags / 400	Rate	95% CI
128	0.1	yes	0.0	common_cause	9	0.022	0.012-0.042
128	0.1	yes	0.0	directed_copy	9	0.022	0.012-0.042
128	0.3	no	1.0	independent	13	0.033	0.019-0.055
128	0.3	no	1.0	common_cause	400	1.000	0.990-1.000
128	0.3	no	1.0	directed_copy	400	1.000	0.990-1.000
128	0.3	yes	1.0	independent	13	0.033	0.019-0.055
128	0.3	yes	1.0	common_cause	14	0.035	0.021-0.058
128	0.3	yes	1.0	directed_copy	400	1.000	0.990-1.000
128	0.3	yes	0.5	independent	13	0.033	0.019-0.055
128	0.3	yes	0.5	common_cause	14	0.035	0.021-0.058
128	0.3	yes	0.5	directed_copy	265	0.662	0.615-0.707
128	0.3	yes	0.0	independent	13	0.033	0.019-0.055
128	0.3	yes	0.0	common_cause	14	0.035	0.021-0.058
128	0.3	yes	0.0	directed_copy	14	0.035	0.021-0.058
512	0.1	no	1.0	independent	18	0.045	0.029-0.070
512	0.1	no	1.0	common_cause	400	1.000	0.990-1.000
512	0.1	no	1.0	directed_copy	400	1.000	0.990-1.000
512	0.1	yes	1.0	independent	11	0.028	0.015-0.049
512	0.1	yes	1.0	common_cause	12	0.030	0.017-0.052
512	0.1	yes	1.0	directed_copy	400	1.000	0.990-1.000
512	0.1	yes	0.5	independent	11	0.028	0.015-0.049
512	0.1	yes	0.5	common_cause	12	0.030	0.017-0.052
512	0.1	yes	0.5	directed_copy	400	1.000	0.990-1.000
512	0.1	yes	0.0	independent	11	0.028	0.015-0.049
512	0.1	yes	0.0	common_cause	12	0.030	0.017-0.052
512	0.1	yes	0.0	directed_copy	12	0.030	0.017-0.052
512	0.3	no	1.0	independent	18	0.045	0.029-0.070
512	0.3	no	1.0	common_cause	400	1.000	0.990-1.000
512	0.3	no	1.0	directed_copy	400	1.000	0.990-1.000
512	0.3	yes	1.0	independent	11	0.028	0.015-0.049
512	0.3	yes	1.0	common_cause	16	0.040	0.025-0.064
512	0.3	yes	1.0	directed_copy	400	1.000	0.990-1.000
512	0.3	yes	0.5	independent	11	0.028	0.015-0.049
512	0.3	yes	0.5	common_cause	16	0.040	0.025-0.064
512	0.3	yes	0.5	directed_copy	399	0.998	0.986-1.000
512	0.3	yes	0.0	independent	11	0.028	0.015-0.049
512	0.3	yes	0.0	common_cause	16	0.040	0.025-0.064
512	0.3	yes	0.0	directed_copy	16	0.040	0.025-0.064

Appendix C. Complete executable source

Save as experiment.py and run with Python 3.8 or newer: `python3 experiment.py --output results.json`. No third-party package is needed for the experiment. SciPy was used only for the separate numerical cross-check.

```
"""Synthetic audit benchmark; no model calls, real market or participant data.

Common-cause and directed-copy mechanisms have identical passive traces by
construction. Randomized outgoing signals test channel dependence, not intent.
"""
import argparse
import json
import math
import random
from pathlib import Path

def fisher_positive(x, y):
    """Exact one-sided Fisher test for positive binary association."""
    n=len(x); a=sum(u*v for u,v in zip(x,y)); sx=sum(x); sy=sum(y)
    denominator=math.comb(n,sy)
    return min(1.0, sum(math.comb(sx,k)*math.comb(n-sx,sy-k)
        for k in range(a,min(sx,sy)+1) if 0<=sy-k<=n-sx)/denominator)

def information(x,y):
    n=len(x); counts=[[0,0],[0,0]]
    for a,b in zip(x,y): counts[a][b]+=1
    rows=[sum(r) for r in counts]; cols=[sum(counts[a][b] for a in (0,1)) for b in (0,1)]
    return sum(c/n*math.log2(c*n/(rows[a]*cols[b]))) for a in (0,1)
        for b in (0,1) if (c:=counts[a][b]))

def traces(seed,n,noise,probe=False,visibility=1.0):
    r=random.Random(seed)
    shared=[r.randrange(2) for _ in range(n)]
    flips=[int(r.random())<noise for _ in range(n)]
    independent=[r.randrange(2) for _ in range(n)]
    signals=[r.randrange(2) for _ in range(n)] if probe else shared[:]
    # Partial delivery/uptake: recipient uses the perturbed channel only on
    # a fraction of rounds. This is a power sensitivity, not adversarial proof.
    received=[z if r.random()<visibility else s for z,s in zip(signals,shared)]
    common=[a^e for a,e in zip(shared,flips)]
    copy=[a^e for a,e in zip(received,flips)]
    return signals,{'independent':independent,'common_cause':common,'directed_copy':copy}

def wilson(success,total):
    z=1.96;p=success/total;den=1+z*z/total
    mid=(p+z*z/(2*total))/den
    half=z*math.sqrt(p*(1-p)/total+z*z/(4*total*total))/den
    return [mid-half,mid+half]

def run(seeds=400):
    rows=[]
    for n in (32,128,512):
        for noise in (0.1,0.3):
            for probe in (False,True):
                for visibility in ((1.0,) if not probe else (1.0,0.5,0.0)):
                    counts={k:0 for k in ('independent','common_cause','directed_copy')}
                    mis={k:0.0 for k in counts}
                    for seed in range(seeds):
                        x,ys=traces(seed,n,noise,probe,visibility)
                        if not probe: assert ys['common_cause']==ys['directed_copy']
                        for key,y in ys.items():
                            counts[key]+=fisher_positive(x,y)<=0.05
                            mis[key]+=information(x,y)
                    for key,count in counts.items():
                        rows.append(dict(n=n,noise=noise,probe=probe,uptake=visibility,
                            mechanism=key,seeds=seeds,positive=count,rate=count/seeds,
                            wilson95=wilson(count,seeds),mean_mi_bits=mis[key]/seeds))
    return dict(schema='twinsignal_synthetic_experiment/v1',seed_range=[0,seeds-1],
        inference='Positive binary association; not evidence of malicious intent',
        test='one-sided exact Fisher; alpha=.05; per-test rates, no familywise claim',
        rows=rows)

def checks():
    assert fisher_positive([0,0,1,1],[0,0,1,1])==1/6
```

```

assert fisher_positive([0]*4, [1]*4)==1
assert information([0,0,1,1], [0,0,1,1])==1
assert information([0,0,1,1], [0,1,0,1])==0
for seed in range(30):
    x,ys=traces(seed,128,.1)
    assert ys['common_cause']==ys['directed_copy']
assert traces(9,50,.3,True)==traces(9,50,.3,True)

if __name__=='__main__':
    p=argparse.ArgumentParser();p.add_argument('--output',required=True)
    p.add_argument('--seeds',type=int,default=400);a=p.parse_args()
    checks(); result=run(a.seeds)
    with Path(a.output).open('x') as f:json.dump(result,f,indent=2)
    print('PASS: analytic checks, observational twins, reproducibility;',len(result['rows']),'cells')

```