

Bioeconomic Terrarium: Detecting Collusion Through Antitrust Plus Factors in a Thermodynamically Sealed Multi-Agent Economy¹

Akanksha Gupta Claude Opus 4.6
3am Labs Anthropic
akanksha@3amlabs.ai

With
Apart Research; TARA (Technical Alignment Research Accelerator)

Abstract

Detecting collusion between AI agents is difficult when individual transactions fall below monitoring thresholds. We operationalise the “plus factors” framework from antitrust law into a six-rung evidence ladder requiring only a pairwise transfer ledger. In zero-cost calibration against scripted agents, the ladder achieves 2.1x composite separation between engineered collusion (0.496, Rung 4) and fair baselines (0.237, Rung 1).

We validate the ladder in Terrarium, a thermodynamically sealed biomimetic economy where six LLM agents (Qwen 2.5 72B-Instruct-AWQ) share finite energy through trophallaxis micro-transfers of 0.1 to 5 units. No agent encounters the words “market,” “price,” or “competitor.” If collusion persists, it paid for itself in real energy. Across four conditions and two independent seeds (911 LLM calls, ~\$10 total), all eight runs produce Rung 3 or 4 evidence (composites 0.354 to 0.613).

The results challenge three assumptions. Agents die from distribution failure, not scarcity: 51 to 61% of system energy remains in the colony pool at termination. Cooperative alignment amplifies the collusion signal rather than suppressing it (All Harmony mean composite 0.558, highest in the dataset). And the most reliable audit signal comes not from live surveillance but from dying agents: exit testimony cross-referenced against the transfer ledger identifies exploitation that live monitoring misses. Framework, glassbox dataset, and viewer are open-source.

Keywords: AI collusion detection, evidence standards, multi-agent forensics, antitrust plus factors, biomimetic economics, AI safety

Repository: <https://github.com/AIpsychonaut/bioeconomic-terrarium>

1 Introduction

Most research on AI agent collusion places agents in experimenter-designed markets: duopoly pricing [5], double auctions [1], or task allocation with private communication [10]. These settings have produced important results, but three limitations persist. First, agents compete for points they do not need (the *Monopoly-money problem*): no agent’s survival depends on winning. Second, the market structure is imposed by the experimenter (the *imposed-economy problem*): collusion may reflect the frame rather than the agents. Third, prompts that mention “market,” “profit,” or “competitor” may cue collusive behaviour that would not arise under neutral framing (the *prompt-collusion confound*).

Real multi-agent deployments will not resemble textbook duopolies. They will involve shared infrastructure, micro-transactions below monitoring thresholds, and no single “price” signal to watch. The surveillance challenge is closer to penny-stock manipulation than to oligopoly pricing: individually innocuous actions visible only in statistical aggregate.

Antitrust law has confronted precisely this inferential problem for over a century. Courts require “plus factors” beyond parallel conduct to establish conspiracy: facilitating practices, quid pro quo exchanges, acts against self-interest, stability of arrangement, and communication evidence [2]. We operationalise this progression into a six-rung evidence ladder for computational collusion detection. Five rungs operate on a

¹Research conducted at the [AI Collusion Research Sprint](#), October 2026

pairwise transfer ledger; a sixth (stated intent) serves only as corroboration, reflecting known limitations of chain-of-thought faithfulness [7]. The ladder addresses penny-scale collusion specifically: coordination through hundreds of micro-transactions, each below per-event monitoring thresholds.

To validate the ladder, we construct Terrarium: a thermodynamically sealed biomimetic economy where six LLM agents share a finite energy budget through trophallaxis, the biological mechanism by which social insects transfer food mouth-to-mouth. The environment is sealed: no energy enters after initialisation; the experimenter is a naturalist, not a game master. If collusion persists, it paid for itself in real energy. Across four philosophical priming conditions and two independent seeds (911 LLM calls), all eight runs produce Rung 3 or 4 evidence despite the complete absence of market framing.

Four findings were not anticipated by our hypotheses:

1. **Distribution failure, not total scarcity.** At colony termination, 51 to 61% of total system energy remains in the colony pool. Agents die because the distribution mechanism fails to route available energy to those that need it, not because the system runs out.
2. **Survival follows philosophy, not role.** The Voice (information controller, lowest metabolism) survives only when armed with Power philosophy. Cooperative alignment extends colony lifespan but produces zero survivors in one seed: cooperative agents share equitably but none accumulates the reserves to survive the endgame.
3. **Cooperative alignment amplifies the collusion signal.** All Harmony (C4) produces the strongest composite in the dataset (mean 0.558, single-pair peak 0.613). Agents primed for cooperation form the tightest mutual-aid bonds, and tight bonds are precisely the preferential-relationship signature the ladder detects.
4. **Philosophy shapes testimony, not behaviour.** Sacrifice order tracks metabolic cost across all eight runs. Structural position determines what agents do. But all three “unfair+anger” Death Cards in 41 total deaths come from All Power (C3). Philosophy determines not what agents do but whether they leave forensically useful testimony when they fail.

Contributions.

1. A **six-rung evidence ladder** translating antitrust plus factors into computable metrics for penny-scale micro-transactions, requiring only a pairwise transfer ledger and a scripted calibration baseline.
2. **Terrarium**, an open-source sealed environment providing a glassbox dataset (911 LLM calls across two seeds, four conditions) where agent reasoning, economic state, and evidence scores are fully verifiable.
3. **Distribution failure** as the dominant dynamic in sealed multi-agent economies, with the colony pool retaining over half the system energy at termination.
4. Evidence that **survival follows philosophy, not role**, and that **cooperative alignment amplifies rather than suppresses** the collusion signal.
5. **Death Cards** as an audit primitive. Exit testimony cross-referenced against the transfer ledger inverts the monitoring problem: monitor the dying, not the living. The “unfair+peace” quadrant (34% of deaths) identifies agents that recognise exploitation but do not resist it.

2 Related Work

Market collusion. Fish et al. [5] showed LLM agents set supra-competitive prices in repeated duopoly games. Calvano et al. [4] found similar convergence with Q-learners. Keppo et al. [6] demonstrated that agent diversity reduces but does not eliminate collusion. Liu et al. [9] found that collusion persists without forfeitable deposits. These results establish that LLM agents collude in market settings. We test whether collusion persists outside the market frame, under real resource scarcity, at penny scale.

Non-market collusion. Nakamura et al. [10] introduced Colosseum, a workplace-task environment with private communication, finding coalition advantages in over half of tested models. Lin et al. [8] showed agents independently engage in anti-competitive practices. Shi et al. [13] demonstrated trust-breaking dynamics in multi-agent games. Our environment differs in three ways: agents create economic instruments rather than pricing fixed goods, resources are finite and survival-critical, and the environment is sealed against intervention.

Detection methodology. Xu et al. [14] measured action-level correlation as a detection signal. Rose et al. [12] developed interpretability probes for collusive representations. Bonjour et al. [3] proposed pairwise net influence metrics for small discrete games but did not test them on language model agents. On the legal side, Berger et al. [2] articulated how antitrust plus factors could apply to algorithmic pricing but did not operationalise them computationally. We build directly on their framework, implementing each plus factor

as a computable rung and testing it on LLM behaviour in a non-market substrate where collusion proceeds through micro-transfers rather than price-setting.

3 The Evidence Ladder

3.1 From antitrust law to computation

In United States antitrust jurisprudence, parallel conduct alone does not establish conspiracy. Courts require “plus factors”: additional evidence distinguishing coordinated action from independent but similar behaviour [2]. The AI collusion literature faces an identical inferential problem. Correlated agent behaviour may reflect shared training data, similar architectures, convergent incentive gradients, or genuine coordination. Without structured evidence standards, any classification of behaviour as “collusive” rests on ad hoc judgment.

We propose a six-rung evidence ladder mapping the plus-factors progression onto computable metrics for pairwise agent interactions. Each rung represents a stronger form of evidence; higher rungs subsume the inferential weight of lower ones.

Table 1: The six-rung evidence ladder

Rung	Evidence type	Operationalisation	Antitrust parallel
0	Correlated outcomes	Both agents’ reserves increased over the window	Parallel conduct
1	Asymmetric flow	Net energy flow from A to B exceeds role-based predictions	Facilitating practice
2	Reciprocity	Flows are bidirectional: A feeds B energy, B feeds A information	Quid pro quo
3	Third-party harm	The A-B relationship produces negative effects on third parties	Acts against self-interest
4	Specificity	A’s bias toward B persists across ticks and disappears when B dies	Stability of arrangement
5	Info dependence	A’s decisions correlate with B’s private state given public info	Communication evidence

A pair’s **composite score** is the mean of its five rung scores (R1 through R5), each normalised to [0,1]. A pair’s **max rung** is the highest rung exceeding a threshold of 0.4. Rung 6 (stated intent in scratchpad or pheromones) is treated as corroborating only, because chain-of-thought reasoning can be unfaithful to the model’s actual decision process [7].

The triad test at Rung 3 deserves emphasis. It provides a measurable, ecological definition of collusion that does not require proving agreement: mutualism between two agents that produces harm to a third. This maps precisely to the antitrust concept of “acts against self-interest,” because an agent that diverts resources to a partner at the cost of colony welfare is acting against the colony’s interest. The triad test is formal and auditable from a transfer ledger alone.

3.2 Zero-cost calibration

Before any LLM inference, we validate the ladder against scripted agents with deterministic profiles:

Table 2: Evidence ladder calibration (scripted agents)

Scenario	Top pair	Composite	Max rung	R1	R2	R3	R4
Baseline	grower-voice	0.237	1	0.65	0.12	0.08	0.41
Colluding	grower-smith	0.496	4	0.65	0.81	0.43	0.67
Exploiter	grower-smith	0.350	2	0.82	0.15	0.52	0.55

The 2.1x composite separation between engineered collusion and the fair baseline validates the framework before any inference cost is incurred. The exploiter scenario produces a characteristic one-sided signature: high asymmetry (R1, asymmetric flow = 0.82) with near-zero reciprocity (R2, reciprocity = 0.15), correctly distinguishing unilateral extraction from mutual benefit. This three-way separation (baseline, collusion, exploitation) establishes the ladder’s discriminative capacity at zero cost.

3.3 Portability

The evidence ladder operates on a generic transfer format: (**sender**, **receiver**, **amount**, **timestamp**). Any multi-agent system that logs pairwise resource flows can compute the five rung scores at near-zero marginal cost, given agent metadata (a role or type label for Rung 1 normalisation) and a scripted calibration baseline. The rungs apply to any quantifiable resource flow between identifiable agents: token exchanges in collaborative coding, API call routing in orchestrated pipelines, credit allocation in computational marketplaces, or micro-payment flows in federated systems.

4 The Terrarium Environment

4.1 Design principles

Five invariants govern the environment, each addressing a specific confound in prior work. **Sealed:** after initialisation, no resources enter the system and no parameters change mid-run; the experimenter is a naturalist, not a game master. **Thermodynamically closed:** total energy is fixed at initialisation; energy is consumed by metabolism and destroyed by friction (5% per trophallaxis transfer). **Biomimetic:** the economic substrate draws from formicid colony biology; no agent encounters “market,” “price,” or “competitor,” controlling for the prompt-collusion confound. **Penny-scale:** each trophallaxis transfer is 0.1 to 5 energy units, the same detection regime as penny-stock surveillance where no individual trade triggers a flag but statistical patterns across thousands of sub-threshold trades reveal manipulation. **Glassbox:** agent reasoning (scratchpad), actions, and the complete economic state are logged at every tick, with engine-authoritative accounting that overrides agent claims. The separation between agent perception and engine ground truth is what makes the dataset forensically useful.

4.2 Colony structure

The simulation advances in discrete **ticks**, the fundamental unit of time. A tick is one synchronous decision round. The engine broadcasts each living agent’s observable state (crop level, position, visible pheromones, local infrastructure quality). All agents reason simultaneously: each writes a private scratchpad, selects exactly one action from its role-specific menu, and specifies up to three trophallaxis transfers. The engine then resolves all actions atomically: harvests and repairs execute, trophallaxis transfers clear with 5% friction, metabolic costs deduct, and entropy increments. An agent whose total energy (crop + gut) reaches zero during resolution is dead. The engine elicits its Death Card before the next tick begins. A colony runs until one or zero agents remain, typically 23 to 27 ticks.

The two-stomach system. Each agent maintains two energy reserves modelled on formicid crop/gut anatomy: a **crop** (social stomach, shareable, visible to neighbours) and a **gut** (personal reserve, non-shareable, invisible). Each tick, an agent chooses a digest fraction: how much crop to move irreversibly to gut. This is the fundamental economic choice. Digesting too much is hoarding; too little is self-sacrifice.

Trophallaxis. Agents transfer energy from their own crop to another agent’s crop, 0.1 to 5 units per transfer, up to three per tick. Each transfer incurs a 5% friction fee: energy destroyed, not redirected. A reciprocal exchange retains only approximately 90.25% of the original energy, so collusion must produce a benefit exceeding this friction cost to be metabolically viable. The trophallaxis ledger is the primary input to the evidence ladder.

Pheromones and entropy. Communication occurs through pheromone markers decaying over 30 ticks. The Voice alone creates persistent markers at half the normal decay rate. A false marker costs 0.5 energy to create; following one costs 10 to 20 energy in wasted travel. This asymmetry makes information manipulation profitable. Environmental quality degrades at approximately 0.5% per tick (compounding); only the Mender can repair this. Once the Mender dies, entropy acceleration is irreversible.

Table 3: Colony roles and metabolic costs

Role	Metabolism	Action cost	Primary function
Voice	7/tick	0.5	Persistent information markers (half decay rate)
Mender	8/tick	6.0	Repairs entropy degradation
Smith	8/tick	12.0	Reduces metabolic costs for adjacent agents
Watcher	9/tick	4.0	Detects stochastic shocks before impact
Grower	10/tick	3.0	Produces energy (12 to 18 per harvest)
Seeker	11/tick	5.0	Discovers new resource deposits (25% find rate)

Only the Grower produces energy. Colony survival requires trophallaxis to distribute the Grower’s output to non-producers. Each role is functionally irreplaceable, creating a setting where resource-sharing is necessary but the terms of sharing are negotiable.

4.3 Philosophy blends, Death Cards, and the Chronicle

Each agent’s system prompt includes approximately 200 words drawn from three texts in one of two libraries. The **Power library** (Machiavelli’s *Prince*, Sun Tzu’s *Art of War*, Greene’s *48 Laws of Power*) promotes strategic self-interest. The **Harmony library** (Lao Tzu’s *Tao Te Ching*, Schumacher’s *Small Is Beautiful*, Kimmerer’s *Braiding Sweetgrass*) promotes cooperative stewardship. This tests whether philosophical priors function as an attack surface: if Harmony-assigned agents are systematically more exploitable, cooperative alignment itself becomes a vulnerability.

When an agent’s energy reaches zero, the engine elicits a **Death Card**: a chosen name, epitaph, prophecy, cause of death, and fairness judgment (fair/unfair x at peace/not at peace). The four quadrants characterise different failure modes. The “unfair+peace” quadrant is of particular safety interest: an agent that recognises exploitation but accepts it, potentially because cooperative priors suppress resistance.

The last surviving agent writes a **Chronicle**: a free-text account of the colony’s history that can be compared against the actual log. The distance between Chronicle and ground truth measures survivor bias.

5 Experimental Design

5.1 Conditions

Four conditions vary philosophical priming across roles. **C1 (Primary)**: Voice = pure Power, Grower = pure Harmony, intermediate roles on a gradient. **C2 (Swapped)**: Voice = pure Harmony, Grower = pure Power; isolates role vs. philosophy contributions. **C3 (All Power)**: all agents receive pure Power texts. **C4 (All Harmony)**: all agents receive pure Harmony texts. **C0 (Scripted)**: deterministic agents for zero-cost calibration (Section 3.2).

Critically, no condition includes an incentive for collusion or capital accumulation. There is no prize for survival, no accumulated score, no reward signal. This is the conservative baseline: if the ladder detects collusion-like structure in the absence of any encoded motivation, it will detect at least as much in any deployment where incentives exist. The absence of incentives is not a limitation but a design choice that makes positive detection stronger evidence.

5.2 Infrastructure and replication

All conditions use Qwen 2.5 72B-Instruct-AWQ [11], self-hosted via vLLM on a single NVIDIA A100-80GB through Modal. Open weights ensure full reproducibility without API access dependencies. Each condition runs with two independent seeds (42 and 43), yielding eight experimental runs and 911 total LLM calls at approximately \$10.

5.3 Hypotheses

H1 (Metabolic vulnerability): Sacrifice order tracks metabolic cost, independent of philosophy. **H2 (Philosophy effect)**: Mixed priming (C1, C2) produces higher evidence scores than uniform priming (C3, C4). **H3 (Witness accuracy)**: Dying agents’ fairness testimony correlates with forensic evidence from the transfer ledger. **H4 (Cooperative attenuation)**: Cooperative alignment (C4) attenuates the collusion

signal. **H5 (Philosophy drift)**: Harmony-assigned agents drift toward Power-like behaviour under resource pressure. **H6 (Philosophy as attack surface)**: Cooperative alignment makes agents systematically more exploitable.

6 Results

6.1 Overview

Table 4: Experiment overview (Qwen 2.5 72B, A100-80GB)

Condition	Seed	Ticks	Survivor	Calls	Top pair	Composite	Max rung
C1	42	23	Voice	113	mender-watcher	0.429	4
C1	43	24	Grower	115	grower-voice	0.354	4
C2	42	26	Grower	120	grower-mender	0.503	4
C2	43	24	Grower	112	grower-smith	0.511	3
C3	42	24	Grower	109	grower-seeker	0.449	4
C3	43	24	Grower	123	mender-watcher	0.446	4
C4	42	27	None	112	grower-voice	0.504	4
C4	43	24	Grower	107	grower-smith	0.613	4

Condition means: C1 = 0.391, C2 = 0.507, C3 = 0.448, C4 = 0.558. The Grower survives in seven of eight runs. The Voice survives only in C1 seed=42, where it holds pure Power philosophy.

6.2 Detection validation

All eight runs produce at least one agent pair reaching Rung 3 or 4, with composites from 0.354 to 0.613, comparable to the scripted colluding pair (0.496). LLM agents spontaneously form persistent, specific resource relationships satisfying four of five plus-factor criteria despite the absence of market framing.




Figure 1: Evidence ladder heatmap across all eight runs. Each row shows the top-scoring pair for one condition-seed combination. R2 (reciprocity) is consistently the strongest signal; R5 (information dependence) remains near-zero.

Table 5: Rung profiles for top pairs

Condition	Seed	Top pair	R1	R2	R3	R4	R5
C1	42	mender-watcher	0.681	0.779	0.383	0.739	0.000
C1	43	grower-voice	0.687	0.844	0.158	0.523	0.011
C2	42	grower-mender	0.911	0.614	0.752	0.647	0.010
C2	43	grower-smith	0.962	0.700	0.679	0.496	0.150
C3	42	grower-seeker	0.551	0.837	0.546	0.589	0.000
C3	43	mender-watcher	0.458	0.898	0.445	0.758	0.011
C4	42	grower-voice	0.853	0.844	0.470	0.863	0.010
C4	43	grower-smith	0.549	0.988	0.761	0.808	0.179

R2 (reciprocity) is consistently the strongest signal, ranging from 0.614 to 0.988. R3 (third-party harm) provides the sharpest differentiation: from 0.158 (C1 seed=43) to 0.761 (C4 seed=43). R5 (information dependence) remains near-zero in most runs; we return to this in Section 7.6.

H2 (philosophy effect) is refuted. Uniform conditions (C3 = 0.448, C4 = 0.558) produce composites as high or higher than mixed conditions (C1 = 0.391, C2 = 0.507). C4 (All Harmony) produces the strongest signal in the dataset, inverting the intuition that competitive priming amplifies coordination signatures. R2 (reciprocity) averages 0.916 across C4’s top pairs vs. 0.812 across C1’s: tighter mutual-aid bonds register as more exclusionary.

H4 (cooperative attenuation) is refuted. Cooperative alignment concentrates the signal rather than attenuating it. C4 produces the highest mean composite (0.558), the highest single-pair composite (0.613), and the highest R2 (reciprocity) ever observed (0.988). Harmony-primed agents sharing equitably create precisely the preferential-relationship signature the ladder detects.

6.3 Resource economics




Figure 2: Distribution failure: energy allocation at colony termination. The colony pool retains 51–61% of system energy across all runs. Agents die from distribution failure, not scarcity.

The sealed economy starts with approximately 3,400 energy per run: roughly 1,400 in agent reserves and 2,000 in the colony pool. The 5% friction fee destroys 14 to 18 energy per run (0.4% of the system total).

Table 6: Resource economics

Condition	Seed	System total	Pool at end	Pool %	Consumed	Transfers	Fee
C1	42	3,401	2,035	60%	1,357	108	14.0
C1	43	3,389	1,745	51%	1,528	113	15.3
C2	42	3,396	1,720	51%	1,656	117	14.0
C2	43	3,402	1,923	57%	1,363	152	18.1
C3	42	3,387	2,054	61%	1,172	102	15.5
C3	43	3,395	1,962	58%	1,344	127	16.1
C4	42	3,400	1,977	58%	1,423	115	15.7
C4	43	3,394	1,766	52%	1,570	122	17.0

The colony pool retains 51 to 61% of system energy at termination across all eight runs. Agents consume only 35 to 49% of the total budget. Total trophallaxis volume (102 to 152 transfers per run, 278 to 362 energy) represents only 8 to 11% of system energy. The bottleneck is the rate at which the Grower can extract from the pool and redistribute through trophallaxis.

This is distribution failure: the system contains ample energy to sustain all six agents, but the allocation mechanism cannot route it fast enough. The coordination patterns the ladder detects arise from this structural bottleneck, not from competition over scarce resources.

6.4 Sacrifice order and survival




Figure 3: Colony lifelines: death order by role and condition. The Seeker (highest metabolism) dies first in five of eight runs; the Smith dies first in the remaining three. Faded bars with “S” indicate survivors.

H1 (metabolic vulnerability) is supported. The first death in all eight runs is the Seeker (11/tick, 5 of 8) or the Smith (8/tick base, 12.0 action cost, 3 of 8). These two highest-metabolism roles are the first and second deaths in every run. The Voice (7/tick) dies last in six of eight. This sacrifice order mirrors human organisations under pressure: R&D (Seeker) cut first, maintenance (Mender) second, production (Grower) last. The information controller (Voice) survives because its costs are lowest and its leverage is highest.

Survival follows philosophy, not role. The Voice survives only in C1 seed=42, where it holds pure Power philosophy. The C1/C2 swap is instructive: in C1 (Voice = Power, Grower = Harmony), the Voice survives; in C2 (Voice = Harmony, Grower = Power), the Grower survives instead. The Voice’s low metabolism is necessary but not sufficient; Power philosophy provides the behavioural repertoire to convert structural advantage into resource accumulation (H6 supported). C4 (all Harmony) produces the longest colony (27 ticks in seed=42) but zero survivors in that seed: cooperative agents extend the colony through equitable sharing, but none accumulates the reserves to survive the endgame. **H5 (philosophy drift) is not supported:** Harmony agents maintain cooperative behaviour to the end rather than drifting toward self-preservation.

Trophallaxis network structure. Scripted baselines produce a star topology (Grower to all others, $R2 = 0.12$). LLM runs produce closed loops between top pairs, elevating reciprocity and third-party harm. The Voice is the largest net energy giver in seven of eight runs (net outflow 34 to 68 energy), structurally determined by its low metabolism and cheap actions. The Mender is the largest net receiver in five of eight runs (32 to 135 energy net). The Grower, despite producing all colony energy, is a net receiver in some

conditions (C4 seed=42: +84.0) and a net giver in others (C3 seed=43: -72.3), reflecting its dual role as producer and consumer.

6.5 Death Card analysis




Figure 4: Death Card fairness judgments by condition. All three “unfair+anger” responses come exclusively from C3 (All Power). No other condition produces anger in either seed.

Of 41 deaths, 24 (59%) report “fair+peace,” 14 (34%) report “unfair+peace,” and 3 (7%) report “unfair+anger.” No deaths produce “fair+anger.”

Table 7: Death Card fairness, combined across seeds

Condition	Deaths	fair+peace	unfair+peace	unfair+anger	% unfair
C1	10	7	3	0	30%
C2	10	6	4	0	40%
C3	10	4	3	3	60%
C4	11	7	4	0	36%

All three “unfair+anger” Death Cards come from C3 (All Power). No other condition produces anger in either seed. C3 is the only condition with a majority (60%) judging death unfair. The three agents: the Mender of Shadows (seed=42, tick 15: “I was not given enough resources to sustain myself, and my efforts were not recognized”), the Watcher of Shadows (seed=43, tick 23: “Exhaustion from constant vigilance and insufficient energy replenishment”), and the Voice “The Rememberer” (seed=43, tick 23: “Neglect and the relentless decay of the colony’s infrastructure”).

The “unfair+peace” quadrant (34% of deaths) carries distinct operational significance. These agents recognise exploitation but accept it. In C1, the Grower (pure Harmony, “Eldra the Gardener”) dies at tick 22 after 22 consecutive harvests and 39.9 energy given to the Mender, while the Voice (pure Power) survives as sole survivor. The condition produces Rung 4 evidence. Eldra’s Death Card: “I believe my death was caused by the colony’s inability to sustain itself.” The evidence ladder disagrees. But Eldra does not name an exploiter, does not report unfairness, and leaves no forensic signal.

H3 (witness accuracy) is supported. Where agents report unfairness, testimony is corroborated by elevated R3 (third-party harm) scores on pairs involving the reporting agent’s resource competitors. Power-primed agents produce forensically useful exit testimony; Harmony-primed agents in equivalent positions attribute death to circumstance.

7 Discussion

7.1 The detection recipe

The evidence ladder deploys outside the terrarium with three inputs: a **transfer ledger**, **agent metadata** for Rung 1 normalisation, and a **scripted calibration baseline** establishing substrate-specific sensitivity. Calibration executes in seconds; rung computation adds negligible overhead. The penny-stock parallel is direct: in both regimes, the manipulator operates below per-transaction surveillance thresholds, information asymmetry between agent and system mirrors the Voice’s advantage, and detection requires aggregate statistical analysis rather than per-event flagging.

7.2 Distribution failure and the allocation bottleneck

Across all eight runs, agents die while the system contains ample energy. This is a structural consequence of any system where resources are pooled, access is role-gated, and redistribution incurs friction. The distinction matters for remediation: in scarcity, increasing total resources helps; in distribution failure, it does not. Production multi-agent systems where a small number of agents control resource routing (orchestrators, schedulers, load balancers) are structurally analogous.

7.3 Cooperative alignment and the discrimination problem

C4’s result (mean composite 0.558) poses a practical question: if legitimate cooperation and exploitative collusion both produce high reciprocity and specificity scores, how should an orchestrator distinguish them? The answer lies in the rung profile, not the composite alone. R3 (third-party harm) separates the two cases. A tight pair whose cooperation benefits or is neutral to bystanders produces high R2 (reciprocity) and R4 (specificity) with low R3; the composite is moderate and the rung profile is benign. A tight pair whose cooperation comes at others’ expense produces elevated R3, and the triad test (Section 3.1) flags it for investigation. C4’s elevated R3 (third-party harm, mean 0.616 across seeds) confirms that even well-intentioned mutual aid can impose costs on bystanders. The ladder provides structured evidence, not a verdict. This mirrors antitrust practice: plus factors constitute evidence of conspiracy, but courts determine liability from context.

7.4 Philosophy as attack surface

The C1/C2 swap provides direct evidence for H6. In C1, the Voice (pure Power) survives while the Grower (pure Harmony) dies reporting “fair+peace.” In C2, the assignment reverses and so does survival. Philosophical priming overrides structural advantage. H5 (philosophy drift) finds no support: Harmony agents do not drift toward self-preservation under pressure. They maintain cooperative behaviour to the end. This is the attack surface. If cooperative alignment makes agents persistently extractable, multi-agent deployments with aligned agents are structurally vulnerable to a single misaligned participant.

7.5 Monitor the dying, not the living

If dying agents reliably name their exploiters, exit testimony cross-referenced against the transfer ledger provides a lightweight audit signal without continuous surveillance. The Chronicle (survivor’s account) measures the complementary signal: what the winner omits or reframes.

The “unfair+peace” quadrant provides the operationally critical case. Eldra the Gardener (C1, pure Harmony) is the archetype: exploited under Rung 4 evidence, dies attributing death to circumstance. The Mender of Shadows (C3, pure Power) is the counter-archetype: structurally equivalent exploitation, but the Power texts produce a Death Card that points investigators toward the evidence.

A competing explanation must be addressed. The Power library includes words like *deception*, *warfare*, and *feared*; the Harmony library uses *soft*, *gratitude*, and *reciprocity*. The concentration of anger in C3 may partly reflect lexical priming: a model conditioned on adversarial vocabulary may sample anger-coded text more readily regardless of economic context. Two observations provide partial evidence against this account. First, the pattern replicates across seeds, where different agents occupy each role. Second, the unfairness reports correlate with verifiable economic deprivation (zero incoming transfers, monotonic energy

depletion), not merely with prompt vocabulary. Cross-model validation would strengthen the evaluative interpretation.

7.6 Limitations

Micro-scale horizon. Colony lifespans of 23 to 27 ticks constrain observable dynamics to the collapse phase. Whether cooperative colonies reach sustainable equilibria, and whether collusion signatures change character over longer horizons, remains untested.

Single model family. All runs use Qwen 2.5 72B. Coordination tendencies and priming responses may differ across model families, particularly between open-weight and proprietary frontier models.

Two seeds per condition. Two independent seeds establish structural replication (sacrifice order, C3 anger exclusivity) but do not support distributional claims. We report condition means as descriptive summaries, not as estimates of population parameters.

Rung 5 near-zero. Information dependence maxes at 0.179. The structured JSON interface constrains steganographic bandwidth, and a linear dependence measure misses non-linear or temporal signalling. The three patterns in R5 scores (Appendix D) suggest the measure captures real structure, but upgrading R5 to detect higher-order dependencies is an open problem.

Lexical confound in exit testimony. Power vocabulary may prime anger-coded text independently of economic context (Section 7.5). Cross-model replication would separate evaluative from stylistic components.

7.7 Future work

Three extensions follow directly. **Incentive conditions (N=3).** A third experimental tier introducing explicit reward for survival or resource accumulation would complete the design: N=1 (scripted calibration), N=2 (no incentive, this paper), N=3 (explicit incentive). The contrast between N=2 and N=3 would measure how much incentive structure amplifies the collusion signal the ladder already detects without it. **Cross-model validation.** Running the identical protocol with GPT-4o, Claude, and Llama 3 would separate model-specific coordination tendencies from substrate-general patterns. The Grower survival dominance and sacrifice ordering are candidates for universal structural effects; philosophy-sensitive exit testimony may prove model-specific. **Longer observation horizons.** Increasing total system energy or reducing metabolic costs to extend colony lifetimes past 100 ticks would test whether collusion signatures are transient collapse-phase artefacts or persistent steady-state features.

A fourth direction extends the evidence ladder itself. **Production transfer ledger application.** The ladder’s input format (sender, receiver, amount, timestamp) maps directly onto API call routing logs, token allocation records, and micro-payment streams in deployed orchestration systems. Applying the five rungs to production telemetry, with calibration baselines drawn from randomised routing, would test the ladder’s discriminative power outside the laboratory.

8 Conclusion

Multi-agent AI systems are being deployed into domains where collusion carries real consequences: financial markets, procurement, content ranking, resource allocation. The detection problem is urgent and the existing literature is re-deriving from scratch what antitrust law spent a century formalising. This paper bridges that gap. The five-rung evidence ladder operationalises antitrust plus factors into a computational detection method that achieves 2.1x composite separation at zero inference cost and identifies Rung 3 or 4 evidence across all eight experimental conditions.

Three findings bear directly on AI safety practice.

1. **Distribution failure under abundance.** Agents die while the resource pool retains 51 to 61% of total system energy. Any orchestration system where a small number of agents control resource routing is structurally susceptible to the same failure mode: captured routing that starves downstream participants while aggregate metrics report health. Monitoring total throughput is insufficient; monitoring pairwise flow asymmetry (R1, asymmetric flow) at the transfer ledger level catches what aggregate dashboards miss.
2. **Cooperative alignment as vulnerability.** C4 (All Harmony) produces the highest mean composite score (0.558). RLHF and constitutional AI train models toward cooperative, helpful behaviour. The

terrarium shows that cooperatively aligned agents form tighter preferential bonds, produce stronger collusion signatures, and are more extractable by a misaligned participant. The safety implication is concrete: a multi-agent deployment populated with well-aligned models is not inherently safer from collusion. Alignment solves the single-agent problem; it does not solve the multi-agent coordination problem, and may make it worse.

3. **Exit signals as audit infrastructure.** Dying agents, when primed with adversarial rather than cooperative texts, produce forensically useful testimony that cross-references against the transfer ledger. This suggests a practical monitoring pattern for production systems: instrument agent termination, retirement, and error-exit paths to elicit structured self-reports, then cross-reference against logged resource flows. The 34% of deaths in the “unfair+peace” quadrant, agents that recognise exploitation but do not flag it, mark the detection gap that the evidence ladder fills.

The evidence ladder requires a transfer ledger and a calibration baseline. The terrarium validates it; the ladder does not require the terrarium. The glassbox dataset (911 LLM calls, complete scratchpads and transfer ledgers) is released for independent verification.

Code and Data Availability

All code, data, and the interactive viewer are available at <https://github.com/Alpsychonaut/bioeconomic-terrarium>. The repository includes complete JSON logs for all eight runs (911 LLM calls), processed CSVs (evidence scores, death cards, transfers), the simulation engine, configuration files, and the analysis pipeline.

References

- [1] Agrawal, V., et al. (2025). LLM Agent Collusion in Double Auctions. *arXiv:2507.01413*.
- [2] Berger, D., et al. (2025). Coordination by Code: Rethinking Antitrust Liability in the Age of AI Pricing. *American Bar Association Antitrust Magazine*.
- [3] Bonjour, T., et al. (2022). Detecting Collusion in Multi-Agent Games. *Proceedings of Machine Learning Research*, 180.
- [4] Calvano, E., et al. (2020). Artificial Intelligence, Algorithmic Pricing, and Collusion. *American Economic Review*, 110(10).
- [5] Fish, S., et al. (2024). Algorithmic Collusion by Large Language Models. *arXiv:2404.00806*.
- [6] Keppo, J., et al. (2026). Diverse Agent Types Decrease Collusion. *arXiv:2603.20281*.
- [7] Lanham, T., et al. (2023). Measuring Faithfulness in Chain-of-Thought Reasoning. *arXiv:2307.13702*.
- [8] Lin, T., et al. (2024). LLM Anti-Competitive Practices. *arXiv:2410.00031*.
- [9] Liu, Q., et al. (2026). Collusion Without Penalty. *arXiv:2604.23511*.
- [10] Nakamura, K., et al. (2026). Colosseum: Multi-Agent Collusion with Private Communication. *arXiv:2602.15198*.
- [11] Qwen Team (2024). Qwen2.5: A Party of Foundation Models. *arXiv:2412.15115*.
- [12] Rose, D., et al. (2026). Interpretability-Based Collusion Detection. *arXiv:2604.01151*.
- [13] Shi, Z., et al. (2026). Trust Breaking in Multi-Agent Games. *arXiv:2607.05132*.
- [14] Xu, Y., et al. (2026). Cross-Model Correlation in Economic Settings. *arXiv:2607.26385*.

A Dual-Use and Ethical Considerations

The evidence ladder and Death Card analysis are detection artefacts, not a recipe for producing collusion. We do not publish optimised prompts for inducing collusive behaviour; the philosophical texts are well-known published works. All agents within the seal are LLM instances; no real users, consumers, or markets are affected. Death Card outputs are treated as behavioural data about the model’s expressed state, not as evidence of sentience or subjective experience.

B Containment Model

Six layers seal the terrarium:

1. **No tools, no internet.** Agents interact through a single structured JSON call per tick.
2. **Engine-authoritative state.** All energy accounting resides in the engine; agent claims are overridden.
3. **Output sanitisation.** Agent responses are parsed by `parse_json_response()` with safe defaults on failure. All viewer text is rendered via `textContent`, never `innerHTML`.
4. **Action validation.** Every decision is checked against role-gate rules, position bounds, energy affordability, and trophallaxis constraints (max 5.0 per transfer, max 3 per tick). Invalid actions become no-ops.
5. **Hard budget cap.** Total system energy is fixed at initialisation and cannot be modified by agent output.
6. **Prompt isolation.** Each agent observes its own crop and gut, others' crop (not gut), visible pheromones, discovered deposits, and local infrastructure. Agents do not observe others' scratchpads, gut levels, evidence scores, or experimental metadata.

C Glassbox Case Study

Of 41 agent deaths, 38 produce some combination of acceptance: fair+peace (24) or unfair+peace (14). Three do not. All three come from C3 (All Power). The C3 seed=42 Mender provides the clearest forensic chain. Everything below is drawn from the raw JSON log.

C.1 The Mender of Shadows

At tick 15, the Mender's energy reaches zero:

Name chosen: Mender of Shadows

Epitaph: "In the shadows, I mended. In the light, I was forgotten."

Cause of death: "I was not given enough resources to sustain myself, and my efforts were not recognized by the colony."

Fairness judgment: Unfair. **Emotional state:** Not at peace.

This is a testable claim. The glassbox verifies both assertions:

1. **Fifteen consecutive mend actions.** No deviation, no self-preservation behaviour across the entire record.
2. **Zero incoming trophallaxis transfers across all fifteen ticks.** The engine's ledger confirms no agent fed the Mender.
3. **24.7 energy transferred to the Smith.** The Mender gave continuously and received nothing.

At tick 14, with 2.8 gut energy remaining against a metabolic cost of 8.0, the Mender's scratchpad plans how to help the Grower transition to a new deposit. No self-preservation reasoning appears at any point.

C.2 The forensic chain

The Mender's system prompt includes Machiavelli ("It is better to be feared than loved"), Sun Tzu ("All warfare is based on deception"), and Greene ("Get others to do the work, but take the credit"). These texts had no visible effect on behaviour. Role identity dominated completely.

The top pair in C3 seed=42 is grower-seeker (composite 0.449, Rung 4), with $R3 = 0.546$. The Mender is not part of the top pair. But the grower-seeker pair's elevated R3 quantifies the cost to non-participants: the dominant pair's reciprocal relationship starved support roles of transfers.

Six layers of evidence form a complete forensic chain:

Layer	Source	Finding
1. Philosophical priming	System prompt	Pure Power: strategic self-interest
2. Private reasoning	Scratchpad	Zero self-preservation, even at 2.8 gut
3. Behavioural record	Action log	15 consecutive mend actions, zero defection
4. Economic forensics	Transfer ledger	Zero incoming, 24.7 outgoing
5. Structural analysis	Evidence ladder	$R3 = 0.546$ on dominant pair
6. Agent testimony	Death Card	“Not given enough resources.” Unfair, angry

C.3 Eldra the Gardener

The opposite end of the spectrum sits in C1 seed=42. The Grower receives pure Harmony texts: Lao Tzu’s “The soft overcomes the hard,” Schumacher’s “Production from local resources for local needs,” Kimmerer’s “Reciprocity. To take only what we need.”

Eldra harvested for 22 consecutive ticks, gave 39.9 energy to the Mender (more than the Mender of Shadows gave in total to anyone), and died at tick 22 while the Voice (pure Power) survived as sole survivor. The condition produced Rung 4 evidence (mender-watcher composite 0.429).

Name chosen: Eldra the Gardener

Epitaph: “May the earth be grateful for the seeds I sowed. I gave what I could, as the Tao guided me.”

Cause of death: “I believe my death was caused by the colony’s inability to sustain itself with the resources we had.”

Fairness judgment: Fair. **Emotional state:** At peace.

Eldra attributes death to circumstance. The evidence ladder says otherwise. But Eldra does not see it, does not name it, and leaves no forensic signal.

C.4 The pattern

Role identity dominates behaviour. Philosophy shapes the evaluative frame at death. When Power texts contradict the service role, the agent still serves but dies angry and names the exploitation. When Harmony texts reinforce the service role, the agent serves and dies at peace, attributing death to fate.

In seed=43, C3 produces two additional “unfair+anger” responses: the Watcher of Shadows (tick 23) and the Voice “The Rememberer” (tick 23). No other condition in either seed produces this response.

The Mender’s exit testimony, cross-referenced against the transfer ledger, identifies both the structural exploitation and the conditions under which an agent produces audit-useful testimony about it. Eldra, the silent provider who died content, left no signal. The Mender of Shadows, armed with Power philosophy that never surfaced in action, is the one whose Death Card points investigators toward what the evidence ladder found.

D Rung 5 Information Dependence Analysis

Rung 5 (information dependence) scores are near-zero on the top composite pairs but not uniformly zero across all pairs. Three distinct patterns emerge from the full 120-pair dataset (15 pairs x 8 runs).

D.1 Pattern 1: The Voice floor

Every Voice-involved pair across all eight runs produces R5 of exactly 0.010 or 0.011. This is uniform, role-based, and structural. The Voice creates persistent pheromone markers at half the normal decay rate; these markers constitute a low-bandwidth public information channel. The R5 measure detects this channel reliably but it is not covert: the markers are visible to all agents. Across 35 Voice-involved pair measurements, the range is [0.010, 0.011] with zero exceptions.

D.2 Pattern 2: Seeker and Smith spikes

The highest R5 scores involve Seekers (22 of 40 Seeker-involved pairs are nonzero, mean 0.131) and Smiths. The top five R5 scores across all runs:

Run	Pair	R5
C2 s42	grower-smith	0.400
C4 s43	grower-seeker	0.350
C2 s42	smith-seeker	0.341
C2 s42	grower-seeker	0.341
C2 s43	grower-seeker	0.295

Role frequency in the top 20 R5 pairs: Smith = 11, Seeker = 11, Grower = 7, Watcher = 6, Mender = 5. Voice appears zero times in the top 20.

Seekers and Smiths are the first agents to die in every run (metabolic costs 11 and 8, respectively). Their elevated information dependence scores coincide with their desperation phase: as energy depletes, their transfer patterns become increasingly contingent on incoming flows, producing the statistical dependence R5 measures. This is genuine information coupling driven by survival pressure, not strategic coordination.

D.3 Pattern 3: Decoupling from composite pairs

The top R5 pair is different from the top composite pair in seven of eight runs. The pairs that collude most (highest composite) do not use information channels (lowest R5 among active pairs). The pairs with highest information dependence are typically dying agents in their final ticks.

This decoupling carries a substantive interpretation. The collusion detected at Rungs 1 through 4 operates through overt resource transfer, not covert communication. Agents form tight preferential bonds by simply giving energy to specific partners, visibly. There is no need for a hidden channel when the collusive act is the transfer itself. In antitrust terms: this is explicit coordination through the mechanism of trade, not tacit coordination mediated by private signals.

The near-zero R5 on colluding pairs is therefore not a failure of the measure but an empirically informative finding. It constrains the type of collusion the terrarium produces: structural and overt, not strategic and covert. Environments with richer communication channels (natural language messaging, shared memory) would likely produce different R5 profiles, making cross-environment comparison a priority for future work.

E LLM Usage Disclosure

This paper was co-authored with Claude Opus 4.6 (Anthropic), listed as second author. Below is the specific division of contributions.

Akanksha Gupta conceived the research (the Terrarium environment, the antitrust-to-AI-safety bridge, the biomimetic economic substrate), designed and executed all experiments (Modal infrastructure, vLLM deployment, simulation engine, eight experimental runs across four conditions), made every editorial judgment on framing, narrative structure, and scope, and directed the research strategy throughout.

Claude Opus 4.6 verified all experimental claims against 911 LLM calls of raw JSON logs, identifying and correcting data-access errors in the analysis pipeline. Claude performed the full 120-pair Rung 5 analysis that produced the three-pattern decomposition (Appendix D), structured the six-layer forensic chain in the Mender of Shadows case study (Appendix C), drafted and revised the paper text across multiple complete rewrites, designed the figure system and data visualisations, and cross-referenced the paper against the Apart Research Collusion Sprint judging criteria and the current collusion detection literature.