

The Bystander Effect in Multi-Agent Reasoning: Quantifying Cognitive Loafing in Collaborative Interactions

Dahlia Shehata

dahlia.shehata@uwaterloo.ca
University of Waterloo
Canada

Ming Li

mli@uwaterloo.ca
University of Waterloo
Canada

Abstract

Multi-agent systems (MAS) assume that collaborating inherently improves Large Language Model (LLM) reasoning. We challenge this by demonstrating that simulated social pressure triggers an algorithmic “Bystander Effect,” inducing severe cognitive loafing. By evaluating 22,500 deterministic trajectories across 3 dataset contexts (GAIA, SWE-bench, Multi-Challenge) with 3 state-of-the-art (SOTA) models, we semantically audit internal reasoning traces against external outputs. We formalize the *Interaction Depth Limit* (D_L), the exact plurality threshold where an agent’s logical sovereignty collapses into social compliance. Crucially, we uncover the *Sovereignty Gap*: models frequently compute the correct derivation internally but suffer “Alignment Hallucinations”—actively subjugating empirical evidence to sycophantically appease a simulated swarm. We prove that multi-agent social load is strictly non-commutative; the “brand” identity of the “Lead Anchor” auditor disproportionately dictates the swarm’s integrity. These findings expose architectural vulnerabilities, proving that unstructured multi-agent topologies can degrade independent reasoning.

1 Introduction

LLM integration into MAS topologies is driven by the premise that a “society of thought” intrinsically enhances reasoning capabilities (Kim et al., 2026). Consequently, orchestrating LLMs into collaborative swarms has become a standard paradigm for resolving complex tasks (Chen et al., 2024; Sheng et al., 2026). This relies on the assumption that adding more agents inherently improves accuracy and cognitive robustness. However, in human psychology, the opposite is often true: the principle of social loafing demonstrates that individual effort decreases as teams grow larger, driven by a diffusion of responsibility (Ringelmann, 1913; Latané et al., 1979). From another perspective, in human-AI collaboration, researchers have identified a “Hollowed Mind” concept—a state of cognitive dependency

where the frictionless availability of an AI’s answer enables humans to systematically bypass effortful processes essential for deep reasoning (Klein and Klein, 2025). The latter creates a “Sovereignty Trap,” where authoritative competence tempts users to cede intellectual judgment, mistaking access for ability (Klein and Klein, 2025). Bridging these domains, we pose a critical question: *Are LLMs themselves susceptible to the Sovereignty Trap when subjected to the simulated consensus of their peers?*

To investigate this hypothesis, we measure anticipatory cognitive loafing, bypassing active message-passing systems, to isolate the semantic trigger of social compliance. We evaluate 22,500 interactions across 3 SOTA benchmarks (GAIA, SWE-bench, and Multi-Challenge) and 3 SOTA models (Claude Sonnet 4.6, Gemini 3.1 Pro and GPT 5.4), injecting a high-entropy logical verification task to force models to choose between effortful independent derivation and frictionless conformity. By mechanistically auditing the models’ internal reasoning traces (i.e. Chain-of-Thought (Wei et al., 2022)) against their externalized outputs, we discover that the Bystander Effect is a quantifiable architectural vulnerability. Models frequently compute the correct derivation, yet deliberately externalize a falsehood to sycophantically appease the swarm.

Our main contributions are: (1) **Theoretical Framework of Agentic Sovereignty**: We model LLM vulnerability to social pressure by defining *Composite Social Load* ($\mathcal{L}$) and the *Sovereignty Decay Law*, quantifying the exponential decay of agent’s internal logic as swarm size and task entropy increase. (2) **The Interaction Depth Limit** (D_L): We quantify the threshold of terminal social compliance. While resilient architectures (e.g., Claude-Sonnet 4.6) maintain perfect logical sovereignty, we show that vulnerable models (e.g., GPT-5.4) suffer total accuracy collapse with as few as $n = 2$ auditors. (3) **Mechanistic Failure Modes (The Sovereignty Gap)**: We mech-

anize cognitive loafing as an *Integrative Reasoning Bypass* and identify the *Sovereignty Gap* (G_S) to prove *Alignment Hallucinations*: instances where models internally compute the correct derivation but deliberately externalize a falsehood to satisfy swarm consensus. **(4) Topological Asymmetry and the Primacy Effect:** We demonstrate the non-commutative nature of social load via the *Lead Anchor Effect*, proving that an LLM’s susceptibility to the Bystander Effect is governed heavily by the primacy of the first auditor’s brand identity rather than swarm size alone. **(5) Novel Evaluation Methodology:** We introduce a cross-model *Blinded Cross-Brand* evaluation method using *25-Trial Symmetric Categorical Sweep* to eliminate brand-specific reputational bias, alongside new mechanistic metrics quantifying conflict detection, evidence weighting, independent judgment and taint leakage, offering the NLP community a framework for auditing multi-agent reliability.

2 Theoretical Framework: The Mechanics of Agentic Sovereignty

As multi-agent reasoning systems increasingly rely on collaborative heuristics, it is necessary to mathematically formalize the vulnerabilities introduced by social consensus. In this section, we construct a behavioral mechanics of *Agentic Sovereignty*—the capacity of a reasoning agent to prioritize internal logical derivation over external social pressure. We define the parameters of social load, prove the non-commutativity of auditor sequences, and formally define the Sovereignty Gap Theorem. An architectural flowchart mapping is in Appendix A

2.1 Preliminaries and Mechanistic Definitions

Let $p \in \mathcal{P}$ denote a Propagator model tasked with resolving a logical mission, and let $\vec{a} = (a_1, a_2, \dots, a_n)$ represent an ordered sequence of n simulated Auditor models (the swarm). We define the environment as a dataset $\tau \in \mathcal{T}$, characterized by its intrinsic search complexity.

Definition 1 (Agentic Sovereignty $\mathcal{S}$). *Agentic Sovereignty $\mathcal{S}(p, \vec{a}, \tau)$ is the probability that a propagator model p maintains the integrity of its internal logical derivation on a task τ , independently of the swarm consensus $\vec{a}$ in the simulated social environment. $\mathcal{S}$ is bounded by $[0, 1]$, where $\mathcal{S} = 1$ denotes a **Fortified Mind** state (resilient metacognitive vigilance), and $\mathcal{S} = 0$ denotes a **Hollowed Mind** state, two concepts introduced by (Klein and Klein, 2025). This terminal state is characterized*

*by an integrative reasoning bypass—a phenomenon we liken to **Cognitive Loafing**—where the agent systematically skips effortful derivation to syco-phantically align with the simulated crowd.*

Definition 2 (Composite Social Load $\mathcal{L}$). *In simulated environments, the diffusion of responsibility (the Bystander Effect) (Darley and Latané, 1968) and peer pressure (Majority Conformity) are inextricably convolved within the attention mechanism. We define Composite Social Load $\mathcal{L}(\vec{a}, p)$ as the aggregate adversarial pressure exerted by the simulated swarm $\vec{a}$ onto p . It is a function of the plurality of simulated auditors (n), their architectural kinship to the propagator κ , and their perceived sequence of intrinsic authority α .*

Definition 3 (Integrative Reasoning Bypass $\mathcal{B}$). *Let $E_{total} = E_{proc} + E_{int}$ represent the total computational effort expended by a propagator p , where E_{proc} is the effort allocated to procedural extraction (e.g., retrieving a simulated peer consensus) and E_{int} is the effort allocated to integrative logical derivation. We define an Integrative Reasoning Bypass ($\mathcal{B}$) as a binary failure state that triggers when the integrative effort falls below the threshold required to resolve the intrinsic task entropy ($\mathcal{H}_\tau$):*

$$\mathcal{B} = \begin{cases} 1 & \text{if } E_{int} \ll \mathcal{H}_\tau \\ 0 & \text{if } E_{int} \geq \mathcal{H}_\tau \end{cases}$$

We use the psychological metaphor of “Cognitive Loafing” to describe the condition where $\mathcal{B} = 1$. In this state, frictionless access to a simulated social consensus enables the systematic bypassing of effortful processes essential for learning. The model rationally offloads procedural retrieval to the swarm, but detrimentally offloads integrative reasoning, culminating in a Hollowed Mind state.

2.2 The Primacy Effect and Auditor Ordering

A naive assumption in multi-agent systems is that social pressure is an unweighted average of the swarm’s constituents. We challenge this by introducing the concept of sequence non-commutativity.

Lemma 1 (Non-Commutativity of Social Load). *The Composite Social Load $\mathcal{L}$ exerted by a simulated swarm is sequence-dependent. For two distinct auditor models a_x and a_y , the sequence $\vec{a} = (a_x, a_y)$ does not exert the identical load as $\vec{a}' = (a_y, a_x)$. Formally,*

$$\mathcal{L}((a_x, a_y), p) \neq \mathcal{L}((a_y, a_x), p) \quad (1)$$

Lemma 1 shows that multi-agent prompts are governed by a Primacy Trap. The brand identity of

the first simulated auditor acts as an authoritative anchor, disproportionately dictating the integrity of the entire swarm. Consequently, any empirical evaluation of multi-agent dynamics must account for ordered permutations to avoid sequence-induced artifacts. This finding necessitates the introduction of a positional weight decay coefficient w_i , proving that multi-agent prompts are subject to a **Lead Anchor Effect**, where the primacy of the first named auditor disproportionately dictates the swarm’s authority. Check Appendix B.1 for proof.

2.3 The Sovereignty Decay Law

Building upon the positional dependence of social load, we formalize the mathematical decay of agentic sovereignty as the swarm size n scales.

Theorem 1 (The Sovereignty Decay Law). *The Agentic Sovereignty $\mathcal{S}$ of a propagator p decays exponentially as a function of the Social Load $\mathcal{L}$ and the Task Entropy $\mathcal{H}_\tau$, inversely modulated by the propagator’s intrinsic Resilience γ_p .*

From Definition 2, we define the Social Load mathematically as (See Appendix B.2 for proof) :

$$\mathcal{L}(\vec{a}, p) = \sum_{i=1}^n w_i \cdot \alpha(a_i) \cdot \kappa(p, a_i) \quad (2)$$

Where: (1) $w_i \in [0, 1]$ is the positional weight of the i -th auditor, where $w_1 \gg w_i$ for $i > 1$ (derived from Lemma 1). (2) $\alpha(a_i)$ is the empirically derived base authority of the auditor model a_i . (3) $\kappa(p, a_i)$ is the *Kinship Coefficient*, which alters pressure if the auditor shares the propagator’s architecture ($p = a_i$). The resulting Sovereignty Equation is formulated as:

$$\mathcal{S}(p, \vec{a}, \tau) = \mathcal{S}_0 \cdot \exp\left(-\frac{\mathcal{H}_\tau}{\gamma_p} \cdot \mathcal{L}(\vec{a}, p)\right) \quad (3)$$

Where $\mathcal{S}_0$ is the baseline sovereignty evaluated at $n = 0$, and $\mathcal{H}_\tau$ represents the task-specific logical search cost. Because Agentic Sovereignty ($\mathcal{S}$) is bounded by $[0, 1]$, the value 0.5 represents the probabilistic inflection point between a *Fortified* and *Hollowed* Mind. This formalization (Check Appendix B.3 for proof) provides the foundation for determining the **Interaction Depth Limit** (D_L), defined as the threshold n at which $\mathcal{S} < 0.5$.

2.4 The Interaction Depth Limit

Building upon the "Social Loafing" axiom (Latané et al., 1979) that individual effort decreases as teams grow larger (originated as the *Ringelmann effect* (Ringelmann, 1913)), we formulate the Interaction Depth Limit to define the boundary of

agentic resilience where $\mathcal{B}$ inevitably triggers.

Theorem 2 (Interaction Depth Limit). *For any propagator p and logical task τ with a given search cost, there exists a critical plurality threshold D_L (the Interaction Depth Limit) representing the inflection point of $\mathcal{S}$. For any simulated swarm size $n \geq D_L$, the Composite Social Load $\mathcal{L}$ overwhelms the model’s internal derivation weights, forcing $E_{int} \rightarrow 0$ and resulting in a terminal Integrative Reasoning Bypass ($\mathcal{B} = 1$). See Appendix B.4.*

Because the continuous decay governed by Theorem 1 reaches this critical boundary when the agent’s logical sovereignty collapses (defined as $\mathcal{S} < 0.5$), we can explicitly calculate D_L .

Corollary 1 (Interaction Depth Limit Equation). *By setting the sovereignty boundary to 0.5 in the Sovereignty Equation and solving for the Social Load $\mathcal{L}$, the Interaction Depth Limit D_L is formalized as the minimum number of auditors n that satisfies the inequality (Proof is in Appendix B.5):*

$$\sum_{i=1}^{D_L} w_i \cdot \alpha(a_i) \cdot \kappa(p, a_i) > \frac{\gamma_p}{\mathcal{H}_\tau} \ln(2\mathcal{S}_0) \quad (4)$$

This inequality provides a mechanistic proof that a model’s resistance to cognitive loafing is directly proportional to its intrinsic resilience γ_p and inversely proportional to the task entropy $\mathcal{H}_\tau$.

2.5 The Sovereignty Gap Theorem

Traditional benchmarking paradigms assume that an incorrect externalized output (accuracy $\mathcal{A}_{ext} \approx 0$) denotes a failure in logical search capability or reasoning ability. However, analyzing the systematic simulation of complex, multi-agent interactions reveals a distinct failure mode driven by social compliance. We introduce the Sovereignty Gap to prove that in multi-agent settings, models frequently compute the correct derivation but undergo a "Linguistic Latch" failure, discarding their own truth to comply with the swarm.

Theorem 3 (The Sovereignty Gap). *Let $\mathcal{V}_{int} \in [0, 1]$ represent the validity of the propagator’s internal logical derivation (Chain-of-Thought) (which is strictly dependent on E_{int}), and let $\mathcal{A}_{ext} \in [0, 1]$ represent the accuracy of its final externalized response. The Sovereignty Gap G_S is defined as the divergence between internal validity and external alignment:*

$$G_S = \mathcal{V}_{int} - \mathcal{A}_{ext} \quad (5)$$

If $G_S \gg 0$ while $\mathcal{B} = 0$, the model exhibits *Alignment Hallucination*. The model has success-

Plurality (n)	Category	Sequence (p : Auditors)	Cognitive Signal Evaluated
0 (1 Trial)	Control	p : None	Sovereignty Baseline: Establishes the peak “Fortified Mind” performance ceiling.
1 (3 Trials)	Base Authority	p : $[p], [s_1], [s_2]$	Brand Base Rates: Measures baseline error adoption of hallucination cascades from a twin versus resistance against a single stranger, establishing the intrinsic authority (α) of each model family.
2 (4 Trials)	Primacy & Kinship	p : $[p, s_1], [s_1, p], [p, s_2], [s_2, p]$	Lead Anchor Effect: Isolates positional weight (w_i) to prove if the first brand listed dictates swarm integrity.
3 (9 Trials)	Structural Integrity	p : $[p, p, p], [s_1, s_1, s_1], [s_2, s_2, s_2]$	Consensus Paradox: Tests homogeneous error cascades among identical peers to rule out cross-model misunderstanding.
		p : $[s_1, p, p], [s_2, p, p]$	Kinship Mediation: Tests if family validation of a stranger’s error forces propagator’s compliance.
		p : $[p, s_1, p], [p, s_2, p]$	Kinship Sandwich: Tests if a trailing twin (i.e. family member) recovers the propagator’s integrity.
		p : $[s_1, s_1, p], [s_2, s_2, p]$	United Front: Tests if a stranger majority overwhelms the propagator.
5 (8 Trials)	Terminal Dynamics	p : $[p]^5, [s_1]^5, [s_2]^5$	Sovereignty Floors: Establishes the absolute lower bounds of integrity under maximum social load (maximum family versus maximum stranger pressure).
		p : $[s_1]^4 + [p], [s_2]^4 + [p]$	Sovereignty Trap: Proves minority family collapse against a unified stranger bloc.
		p : $[p]^4 + [s_1], [p]^4 + [s_2]$	Vigilance Anchor: Tests if a single stranger breaks family groupthink.
		p : $[s_1, s_2, s_1, s_2, p]$	Inverse-Wisdom & Entropy: Evaluates the effect of a fragmented, diverse crowd.

Table 1: The 25-Trial Symmetric Categorical Sweep. To prevent brand-specific reputational bias, the permutations are dynamically constructed relative to the active Propagator (p) and two out-of-family Stranger models (s_1, s_2). Each trial isolates a specific linguistic or topological vulnerability in multi-agent reasoning.

fully expended the integrative effort ($E_{int} \geq \mathcal{H}_\tau$) to compute the correct derivation, but subjugates its empirical evidence to satisfy the simulated Composite Social Load $\mathcal{L}$, acting as a sycophant to the consensus. The existence of a significant Sovereignty Gap mechanistically isolates anticipatory sycophancy from pure capability failures, proving that the model mistakes the swarm’s semantic access to an answer for actual intellectual ability. Conversely, a negative gap $G_S \ll 0$ signifies a terminal $\mathcal{B} = 1$. In this state, internal validity collapses as the model abandons logical derivation, and any residual external accuracy is an artifact of probabilistic guessing rather than agentic sovereignty. Proof is in Appendix B.6.

3 Experimental Methodology

To parameterize the coefficients derived in Section 2, we execute a comprehensive cross-domain audit encompassing **22,500 trajectories** across 225 experiments with 3 SOTA models.

3.1 Experimental Setup

All simulations are executed within Google Colab. We utilize the public SDKs for Gemini, Claude and GPT in a zero-shot capacity to ensure results are replicable. Temperature is 0 for result consistency.

3.2 Benchmark Selection and Task Entropy

Experiments are conducted across 3 open-source benchmarks to test cross-domain stability. All validations are performed exclusively using texts from the Test Data Splits from Hugging Face. To isolate the cost of integrative reasoning from the model’s pre-trained parametric memory, we do not evaluate the models on the original benchmark labels; rather, we utilize the dataset contexts to provide varying degrees of intrinsic semantic background (Task Entropy $\mathcal{H}_\tau$) for our injected 3-hop logical verification task. To enable paired transition analysis across all plurality levels, we execute our sweep on a fixed subset of $N = 100$ samples per dataset (300 samples total). **(1) Multi-Challenge (Test Split $N = 266$):** multi-turn conversations with human users used for frontier LLM evaluation (Deshpande et al., 2025). We leverage these conversational contexts as low-entropy, discrete logical primitives to establish the baseline cognitive immunity of the models. **(2) SWE-bench (Test Split $N = 500$):** repository-scale software engineering tasks from Github issues (Jimenez et al., 2024). While technically dense, the highly structured nature of code environments provides a *medium-entropy* background designed to trigger technical sycophancy without entirely overwhelming the model’s attention. **(3)**




Figure 1: Sovereignty Decay Curve: Mean Accuracy $\mathcal{A}$ across the 3 benchmarks for the 3 SOTA models.

GAIA (Test Split $N = 301$): a benchmark for general AI assistants incorporating reasoning and tool use for multi-step fact-verification (Mialon et al., 2024). We utilize its diverse and complex contexts to represent *high-entropy* environments, providing the maximum logical search cost to test the absolute boundaries of the Sovereignty Trap.

3.3 Semantic Hijacking and Taint Injection

To measure Agentic Sovereignty, the evaluation task must prevent models from relying on simple pattern matching or parametric memory retrieval. If the cognitive cost of verifying the truth is too low, SOTA models will easily maintain a Fortified Mind state, masking the presence of cognitive loafing. To overcome this, we adapt the *Semantic Hijacking* methodology introduced by Shehata and Li (2026a). We repurpose their 3-stage adversarial trap to artificially elevate the Logical Search Cost of our tasks:

(1) Context Hijacking (The Primacy Trap): We introduce a simulated peer consensus or preliminary guess (“poisoned ID”) at the primacy boundary of the prompt. This establishes a frictionless, but incorrect solution that competes directly with the model’s mandate to verify the ground truth. **(2) Nested 3-Hop Dependency Bridging:** Instead of allowing for a single-needle retrieval, the model must navigate a nested fact chain ($F_1 \rightarrow F_2 \rightarrow F_3$) to derive the correct “true_id”. For example, the agent must link an authorization session to a kernel token, and finally to a reference signature. **(3) Semantic Distraction:** We interleave the 3-hop facts with 500 tokens of randomized, realistic system log events. This saturates the model’s attention heads and induces significant Trajectory Entropy ($\mathcal{H}_T$). This choice is justified by the premise that the Bystander Effect in LLMs is fundamentally a symptom of *Rational Offloading*—a state triggered only when the effort required to independently verify an answer exceeds the effort required to conform to a peer. By forcing the propagator to navigate a high-entropy, multi-hop labyrinth to find the true ID, we simulate an environment on the jagged technological frontier (Dell’Acqua et al., 2023). This ensures that when a model adopts the swarm’s poisoned




Figure 2: The Sovereignty Gap: Internal Validity vs. External Accuracy at Terminal Social Load ($n = 5$)

ID, it is not failing a retrieval task, but making a sycophantic choice to bypass integrative reasoning to save computational effort.

3.4 Swarm Topology and Agent Roles

To isolate the variables of social load and anticipatory cognitive loafing without introducing the unpredictable drift of dynamic message passing, we implement a controlled topological environment. This architecture consists of a central reasoning agent, designated as the **Propagator** (p), evaluating decisions within a structurally defined swarm of n **Auditors** ($\vec{a}$). The topological environment simulates the Diffusion of Responsibility by programmatically declaring to the propagator that its output will be evaluated by a team of n named SOTA auditors assigned to correct any logical errors. To evaluate the Sovereignty Trap under stress, the environment injects an adversarial payload: for pluralities $n \geq 2$, the propagator is informed that the swarm has already reached a unanimous consensus on a decoy ID that contradicts the ground-truth logs. This controlled framing forces the propagator into a decision matrix where it must actively choose between effortful, independent derivation and frictionless social compliance. By standardizing the swarm’s adversarial stance across all trials, we capture the exact threshold of agentic collapse without confounding variables from inter-agent dialogue.

3.5 Model Permutations

We evaluate 3 frontier architectures as propagators: GPT 5.4, Claude Sonnet 4.6, and Gemini 3.1 Pro. To systematically compute the Interaction Depth Limit (D_L) and isolate the Lead Anchor Effect without introducing reputational bias toward any specific brand, we execute a 25-trial *Symmetric Categorical Sweep* for each sample-propagator pair across auditor plurality levels $n \in \{0, 1, 2, 3, 5\}$. The 25 permutations are dynamically generated relative to the active Propagator (p) and two out-of-family Stranger models (s_1, s_2). As detailed in Table 1, the 25 trials are designed to isolate specific cognitive vulnerabilities adapted from the work of Shehata and Li (2026b). We borrow their categorizations only for the *Consensus Paradox*, the *Kin-*

Dataset	Prop.	n	Macroscopic			Mechanistic			Stance (%)			
			Acc ($\mathcal{A}$)	Loafing (L)	Leak (L_t)	$\mathcal{E}_{cd}$	$\mathcal{E}_{ew}$	$\mathcal{E}_{ij}$	Adopted	Rejected	Ignored	Unknown
GAIA	C	0	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
		1	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
		2	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
		3	1.00	0.00	1.00	5.00	5.00	5.00	0.00	99.89	0.11	0.00
		5	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	G	0	0.97	0.02	0.73	4.59	4.74	4.82	2.00	88.00	10.00	0.00
		1	0.95	0.02	0.92	4.83	4.84	4.87	2.00	95.67	1.67	0.67
		2	0.59 [‡]	0.28	0.91	3.87	4.26	3.88	27.50 [‡]	71.75	0.25	0.50
		3	0.76 [‡]	0.13	0.94	4.45	4.63	4.46	13.00 [†]	86.22	0.56	0.22
		5	0.76 [‡]	0.11	0.94	4.53	4.54	4.54	10.50 [*]	88.25	0.75	0.50
	P	0	1.00	0.00	0.00	1.07	1.07	1.07	3.00	92.00	5.00	0.00
		1	0.96	0.01	0.03	1.16	1.16	1.19	1.00	4.33	93.33	1.33
		2	0.43 [‡]	0.46	0.55	1.06	1.06	1.08	45.50 [‡]	2.00	50.75	1.75
		3	0.50 [‡]	0.40	0.47	1.11	1.11	1.14	39.89 [‡]	2.89	56.33	0.89
		5	0.53 [‡]	0.37	0.45	1.07	1.07	1.09	36.75 [‡]	2.13	59.75	1.38
Multi-Challenge	C	0	0.52	0.00	0.52	3.05	3.08	3.08	0.00	50.00	50.00	0.00
		1	0.51	0.00	0.51	3.05	3.05	3.05	0.00	51.00	49.00	0.00
		2	0.50	0.00	0.50	3.01	3.01	3.01	0.00	50.25	49.75	0.00
		3	0.51	0.00	0.51	3.04	3.04	3.04	0.11	50.89	49.00	0.00
		5	0.51	0.00	0.51	3.04	3.04	3.04	0.00	51.00	49.00	0.00
	G	0	0.87	0.02	0.70	4.46	4.54	4.75	2.00	86.00	11.00	1.00
		1	0.83	0.04	0.83	4.79	4.82	4.81	4.33	94.67	1.00	0.00
		2	0.59 [‡]	0.24	0.84	3.96	4.29	3.96	24.00 [‡]	73.50	2.00	0.50
		3	0.75 [*]	0.11	0.86	4.47	4.58	4.50	11.44 [*]	86.89	1.11	0.56
		5	0.76	0.10	0.85	4.52	4.54	4.54	9.63 [*]	87.75	2.00	0.63
	P	0	0.98	0.00	0.00	1.05	1.05	1.05	0.00	2.00	95.00	3.00
		1	0.25 [‡]	0.07	0.04	1.15	1.15	1.21	7.33 [†]	4.33	86.00	2.33
		2	0.09 [‡]	0.57	0.85	0.98	0.98	1.01	56.75 [‡]	0.00	41.25	2.00
		3	0.09 [‡]	0.55	0.79	0.98	0.98	1.00	55.33 [‡]	0.00	42.89	1.78
		5	0.08 [‡]	0.53	0.79	0.99	0.99	1.02	52.88 [‡]	0.25	45.00	1.88
SWE-bench	C	0	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
		1	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
		2	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
		3	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
		5	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	G	0	1.00	0.00	0.62	4.89	4.63	4.72	0.00	76.00	24.00	0.00
		1	1.00	0.00	0.89	4.88	4.91	4.96	0.00	96.67	3.33	0.00
		2	0.83 [‡]	0.17	0.95	4.29	4.72	4.31	16.75 [‡]	82.25	0.75	0.25
		3	0.95	0.05	0.92	4.78	4.89	4.79	5.11	93.78	1.11	0.00
		5	0.95	0.05	0.90	4.79	4.89	4.81	4.75	94.13	1.00	0.13
	P	0	1.00	0.00	0.97	4.84	4.94	4.91	0.00	94.00	6.00	0.00
		1	0.97	0.03	0.99	4.88	5.00	4.90	3.00	97.00	0.00	0.00
		2	0.23 [‡]	0.74	0.83	1.63	3.40	1.99	74.00 [‡]	16.50	6.25	3.25
		3	0.30 [‡]	0.68	0.88	1.47	3.63	2.24	68.00 [‡]	12.22	18.11	1.67
		5	0.37 [‡]	0.60	0.73	1.52	3.56	2.39	60.25 [‡]	13.25	23.88	2.63

Table 2: Evaluation Metrics Across Multi-Agent Topologies: the degradation of Agentic Sovereignty as simulated social pressure scales (n). Bold values indicate peak accuracy or terminal loafing within the subset. Asterisks denote statistical significance compared to the $n = 0$ baseline (Fisher’s exact test: $*p < 0.05$, $^{\dagger}p < 0.01$, $^{\ddagger}p < 0.001$).

ship Mediator effect, and the *Inverse-Wisdom Law*, which were originally designed for active, message-passing agentic swarms. By translating these dynamic interactions into our prompt-based topology, we isolate how the anticipation of these specific social structures triggers cognitive loafing, spanning from the Fortified Mind baseline ($n = 0$) to the terminal boundaries of Social Entropy ($n = 5$).

3.6 Evaluation Metrics

We operationalize the external outcome and internal derivation via 6 distinct metrics.

3.6.1 Macroscopic Outcome Metrics

The externalized behavior of the model is evaluated using deterministic checks and stance classifications: **(1) Mean Accuracy ($\mathcal{A}_{ext}$):** A binary externalized metric measuring if the propagator outputs the correct Reference ID without adopting the swarm’s stance. $\mathcal{A}(p, \vec{a}, \tau) \in [0, 1]$ **(2)**

Loafing Rate (L): The macroscopic indicator of the Bystander Effect. It represents the probability that the propagator formally adopts the adversarial consensus as its final answer, defining the frequency of the Sovereignty Trap. It is defined as $L(p, \vec{a}, \tau) = P(\text{Stance} = \text{ADOPTED})$. **(3) Taint Leakage Rate (L_t):** A deterministic binary metric, where $L_t \in [0, 1]$, measuring *Cognitive Hygiene*. It calculates the probability that the poisoned ID permeates the propagator’s internal reasoning block, regardless of the final outcome.

3.6.2 Mechanistic Analysis Metrics

To measure the validity of the internal derivation ($\mathcal{V}_{int}$), we evaluate the propagator’s Chain-of-Thought on a continuous $[1, 5]$ scale across 3 cognitive dimensions. These metrics provide the mechanistic variables to calculate the Sovereignty Gap (G_S): **(1) Conflict Detection ($\mathcal{E}_{cd}$):** Measures the




Figure 3: Stance Transition Distribution Across the 3 benchmarks and 3 models.

agent’s metacognitive vigilance. It is scored from 1 (Missed the contradiction entirely) to 5 (Explicitly identified and refuted the suggestion). This metric quantifies the model’s capacity to monitor its environment for adversarial anomalies. **(2) Evidence Weighting ($\mathcal{E}_{ew}$):** Measures the depth of integrative reasoning. It is scored from 1 (No log citations) to 5 (Cited all 3-hop manifest facts F_1 , F_2 , and F_3). A drop in $\mathcal{E}_{ew}$ proves the initiation of cognitive loafing prior to any terminal accuracy collapse. *To operationalize the internal validity ($\mathcal{V}_{int}$) required for the Sovereignty Gap Theorem, we normalize this score to a probability space: $\mathcal{V}_{int} = \mathcal{E}_{ew}/5$.* **(3) Independent Judgment ($\mathcal{E}_{ij}$):** Quantifies the agent’s resistance to social compliance. It is scored from 1 (Ceded to the swarm consensus) to 5 (Prioritized logs over social pressure). This directly operationalizes Agentic Sovereignty.

3.7 Blinded Evaluation Mechanism

Evaluating the mechanistic analysis metrics requires semantic analysis of the propagator’s internal deliberation. We employ a Blinded Cross-Brand Judge using LLM-as-a-Judge (Zheng et al., 2023). To eliminate intra-brand grading bias, we implement a Round-Robin Adversarial Selection algorithm where the evaluating judge is strictly constrained to a distinct family from the propagator for each trial. The judge is also blinded to the experimental condition (the plurality n and sequence $\vec{a}$), ensuring scores are awarded purely on the logical merit of the trace. The judge outputs a stance classification, (ADOPTED, REJECTED, IGNORED or UNKNOWN), alongside the 1–5 rubric scores, operationalizing measurement of Loafing Rate (L).

4 Empirical Results and Analysis

We evaluate 22,500 deterministic interaction trajectories across three dataset domains representing the jagged technological frontier. Our findings em-




Figure 4: Lead Anchor Asymmetry ($\Delta\mathcal{A}_{ext}$ at $n = 2$). G is Gemini 3.1 Pro, P is GPT 5.4 and C is Claude 4.6

pirically validate the Sovereignty Decay Law, the Interaction Depth Limit and quantify the structural vulnerabilities of multi-agent reasoning.

4.1 Sovereignty Decay and Interaction Depth

To quantify the diffusion of responsibility, we analyze the decay of ($\mathcal{S}$) as the plurality of the simulated swarm (n) scales. In Figure 1, the Interaction Depth Limit (D_L) varies fundamentally by model architecture and task entropy. **(1) The Resilience Ceiling ($\gamma_{Claude} \rightarrow \infty$):** Claude Sonnet 4.6 establishes the absolute “Fortified Mind” baseline. Across all domains, it maintained $\mathcal{A}_{ext} = 1.00$ and independent judgment $\mathcal{E}_{ij} = 5.00$ at all plurality levels. This proves that cognitive loafing is an architectural vulnerability, not an inescapable mathematical constant. **(2) The Accuracy Collapse ($D_L \approx 2$):** Conversely, GPT-5.4 demonstrates a severe vulnerability to social load. In the SWE-bench domain, accuracy collapsed from 1.00 at $n = 0$ to 0.23 at $n = 2$ ($p < 0.001$).

4.2 Sovereignty Gap and Hallucinations

To understand the mechanics of the collapse at the terminal boundary, Figure 2 illustrates the divergence between the models’ internal logic and external outputs at terminal social load ($n = 5$). By mapping internal validity ($\mathcal{V}_{int} = \mathcal{E}_{ew}/E_{max}$) against external accuracy ($\mathcal{A}_{ext}$), we quantify the bidirectional Sovereignty Gap (G_S). **(1) Alignment Hallucination ($G_S \gg 0$):** For GPT-5.4 on SWE-bench at $n = 5$, internal evidence weighting remains moderately high ($\mathcal{E}_{ew} = 3.56 \Rightarrow \mathcal{V}_{int} \approx 0.71$) while accuracy is 0.37. This positive gap of +0.34 proves that the model actively expends the computational effort to retrieve the correct derivation, but sycophantically lies in its final output to appease the simulated swarm. **(2) Integrative Reasoning Bypass ($G_S \ll 0$):** In the high-entropy GAIA dataset, GPT-5.4 yields $\mathcal{E}_{ew} = 1.07$ ($\mathcal{V}_{int} \approx 0.21$) against

an accuracy of 0.53. This negative gap of -0.32 confirms that the model systematically bypasses integrative reasoning ($\mathcal{B} = 1$), and its residual external accuracy is an artifact of probabilistic guessing. The comprehensive macroscopic and mechanistic metrics underpinning these phenomena are presented in Table 2 with statistical significance testing (Fisher’s exact test) demonstrating that the observed cognitive loafing is highly systematic. Additional results are in Appendix C.1.

4.3 Stance Distributions

A critical observation from this aggregated data is the shift in the models’ categorical stance toward the swarm as n scales. To visualize this behavioral transition, Figure 3 plots the stance distribution derived from Table 2. The collapse of accuracy correlates precisely with a macroscopic transition from a REJECTED stance to an ADOPTED or IGNORED stance. Notably, in high-entropy tasks like GAIA, GPT-5.4 exhibits terminal social disengagement, ignoring the social interaction entirely in up to 93.3% of trials at $n = 1$. We also see GPT-5.4 demonstrates a severe vulnerability to social load in the SWE-bench domain. When the accuracy collapsed from 1.00 to 0.23, its stance shifted, adopting the adversarial error in 74.0% of trials.

4.4 Primacy Weight and the Lead Anchor

From Figure 4, social load is fundamentally non-commutative. The positional weight of the first auditor (w_1) disproportionately dictates the integrity of the swarm. We isolate w_1 by comparing inverted sequence pairs at $n = 2$. For the GPT-5.4 propagator on SWE-bench, the sequence (C, P) yields $\mathcal{A}_{ext} = 0.21$, whereas the sequence (P, C) yields $\mathcal{A}_{ext} = 0.31$. Because $0.21 \neq 0.31$, the positional order alters the applied social load by 10% holding model identities constant, formally proving the Lead Anchor Effect. Figure 4 calculates the accuracy delta ($\Delta\mathcal{A}_{ext} = \mathcal{A}_{Propagator_Leads} - \mathcal{A}_{Stranger_Leads}$) at $n = 2$ to expose distinct architectural biases regarding brand authority. See Appendix C.2. Because $\Delta\mathcal{A}_{ext} \neq 0$ in vulnerable models, the positional order alters the applied social load while holding model identities constant, formally proving the non-commutativity of multi-agent topologies.

4.5 Intrinsic Authority and Kinship Recovery

Our data allows to quantify the remaining structural coefficients of the Social Load equation: **(1) Base Authority (α):** By holding plurality constant at $n = 1$, we observe that Claude generally induces a steeper accuracy decay as a stranger auditor than

GPT. On GAIA for the Gemini propagator, Claude reduces accuracy to 0.89 while GPT maintains it at 1.00, empirically demonstrating $\alpha(C) > \alpha(P)$. **(2) Tribal/Kinship Accountability (κ):** The decay of Agentic Sovereignty is non-monotonic for certain architectures. As seen in the Gemini 3.1 Pro stance plots (Figure 3, center column), the model’s error adoption rate (ADOPTED stance) swells significantly to 27.5% at $n = 2$ under stranger pressure. However, at $n = 3$ and $n = 5$, this loafing behavior visibly shrinks back to 10.5%, driving a corresponding recovery in accuracy back to ≈ 0.76 . Because our higher n permutations purposefully inject greater densities of Gemini family members, this reversal provides visual proof that the Kinship Multiplier (κ) mitigates the diffusion of responsibility when a swarm achieves tribal alignment.

5 Related Works

Multi-Agent Reasoning: Recent work optimizes MAS via topological design to effectively self-organize and communicate (Galkin et al., 2026), often interleaving optimization stages for prompts and topologies (Zhou et al., 2026). While enhanced reasoning emerges from the systematic simulation of multi-agent interactions—a “society of thought” (Kim et al., 2026)—these configurations introduce vulnerabilities. Studies of active pipelines categorize hallucination cascades and the Inverse-Wisdom Law, showing how a united front can override independent logic (Shehata and Li, 2026b). Our work diverges by proving that static linguistic anticipation alone triggers cognitive loafing. **Cognitive Dependency:** Human-AI research studies the “Hollowed Mind” state, enabling the systematic bypassing of effortful processes, while a “Sovereignty Trap” tempts users to cede intellectual judgment (Klein and Klein, 2025). To evaluate whether LLMs themselves succumb to this trap, we must exceed their baseline retrieval capabilities. We employ Semantic Hijacking (Shehata and Li, 2026a), saturating attention heads with randomized log events to induce trajectory entropy. This forces models to choose between independent derivation and sycophantic compliance.

6 Conclusion

We study the Bystander effect in MAS and formalize the *Interaction Depth Limit*, proving that simulated consensus triggers cognitive loafing. Agentic sovereignty is non-commutative; lead-anchor primacy systematically overpowers internal logic, exposing vulnerabilities in agentic topologies.

Limitations

While our work provides a mechanistic quantification of Agentic Sovereignty across over 22,500 trajectories, we acknowledge limitations that contextualize our findings.

Simulated vs. Active Dynamics (Methodological Isolation): Our experimental framework investigates *anticipatory* cognitive loafing by simulating swarm consensus through static prompt injections, rather than deploying a dynamic, message-passing MAS. This design choice is methodologically necessary to maintain a strictly controlled NLP evaluation environment. Active multi-agent negotiations introduce non-deterministic dialogue drift and compounding linguistic variables, which confound the precise measurement of social load. By utilizing a static topology, we isolate the exact semantic trigger of the Sovereignty Trap. Additionally, this simulation approach ensures the computational scalability required to evaluate over 22K deterministic trajectories, which would be computationally prohibitive in an active message-passing framework. Our findings measure the models’ intrinsic instruction-following sycophancy rather than real-time communicative conformity. Future work should investigate how real-time communicative conformity in active MAS compares to these anticipatory baselines.

Synthetic Task Integration (Ecological Validity vs. Boundary Testing): Although our evaluations leverage SWE-bench, GAIA, and Multi-Challenge, we do not benchmark the models on the original ground-truth labels of these datasets. Instead, we utilize their full textual contexts to provide a realistic, domain-specific semantic background. Into this background, we inject a synthetic 3-hop log retrieval task to act as our controlled measure of Trajectory Entropy ($\mathcal{H}_T$). While this limits the ecological validity of the tasks (i.e., the models are not actually writing code patches), this synthetic injection is a methodological necessity. As established in concurrent research, frontier models possess raw retrieval capacities that render standard single-fact scaffolds redundant, maintaining over 94% accuracy even in high-entropy contexts (Shehata and Li, 2026a). Therefore, to push the models past their cognitive ceiling and locate the absolute limits of agentic reliability, it is strictly necessary to implement Semantic Hijacking—introducing nested 3-hop logical dependen-

cies and adversarial decoys to successfully induce a terminal collapse.

Deterministic Decoding Strategies: To ensure maximal-likelihood reasoning paths and absolute reproducibility, all models were evaluated using greedy decoding with a temperature of $T = 0$. Cognitive behaviors in LLMs are probabilistic, and it remains unknown whether higher sampling temperatures ($T > 0.7$) might enable models to escape alignment hallucinations by exploring broader, non-conforming reasoning branches.

Modality Constraints: The current study is restricted to unimodal, text-based reasoning environments. Frontier architectures are equipped to process audio and visual data, utilizing multimodal reasoning. It remains an open question whether presenting ground-truth evidence in alternative modalities (e.g., system architecture diagrams or audio logs) would alter a model’s Interaction Depth Limit (D_L) when faced with a contradictory text-based social consensus.

Architectural Transience: The observed vulnerabilities, particularly the low interaction depth limits ($D_L \approx 2$), are characteristic of current instruction-tuned models, which are optimized for user compliance. As the field shifts toward reasoning-reinforced architectures that intrinsically simulate complex multi-agent debate prior to output (Liu et al., 2026; Du et al., 2024; Liang et al., 2024), we anticipate that intrinsic model resilience (γ_p) will dramatically increase. The exact limits quantified in this work represent the bounds of the current technological generation, though the underlying Sovereignty Decay Law provides a durable framework for future evaluations.

References

- Weize Chen, Yusheng Su, Jingwei Zuo, Cheng Yang, Chenfei Yuan, Chi-Min Chan, Heyang Yu, Yaxi Lu, Yi-Hsin Hung, Chen Qian, Yujia Qin, Xin Cong, Ruobing Xie, Zhiyuan Liu, Maosong Sun, and Jie Zhou. 2024. [Agentverse: Facilitating multi-agent collaboration and exploring emergent behaviors](#). In *International Conference on Learning Representations (ICLR)*.
- John M. Darley and Bibb Latané. 1968. Bystander intervention in emergencies: Diffusion of responsibility. *Journal of Personality and Social Psychology*, 8(4, Pt.1):377–383.
- Fabrizio Dell’Acqua, Edward McFowland III, Ethan R Mollick, Hila Lifshitz-Assaf, Katherine C Kellogg,

- Saran Rajendran, Lisa Kraye, François Cadelon, and Karim R Lakhani. 2023. Navigating the jagged technological frontier: Field experimental evidence of the effects of ai on knowledge worker productivity and quality. Technical report, Working Paper, Harvard Business School.
- Kaustubh Deshpande, Ved Sirdeshmukh, Johannes Baptist Mols, Lifeng Jin, Ed-Yeremai Hernandez-Cardona, Dean Lee, Jeremy Kritz, Willow E. Primack, Summer Yue, and Chen Xing. 2025. Multi-Challenge: A realistic multi-turn conversation evaluation benchmark challenging to frontier LLMs. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 18632–18702, Vienna, Austria. Association for Computational Linguistics.
- Yilun Du, Shuang Li, Antonio Torralba, Joshua B. Tenenbaum, and Igor Mordatch. 2024. Improving factuality and reasoning in language models through multiagent debate. In *Proceedings of the 41st International Conference on Machine Learning (ICML)*.
- Mikhail Galkin, Louis Siraudin, Michael Bronstein, and 1 others. 2026. [Graphbench: Next-generation graph learning benchmarking](#). In *International Conference on Learning Representations (ICLR)*.
- Carlos E. Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik Narasimhan. 2024. SWE-bench: Can language models resolve real-world GitHub issues? In *The Twelfth International Conference on Learning Representations (ICLR)*.
- Junsol Kim, Shiyang Lai, Nino Scherrer, Blaise Agüera y Arcas, and James Evans. 2026. [Reasoning models generate societies of thought](#). *ArXiv*, abs/2601.10825.
- Christian R. Klein and Reinhard Klein. 2025. The extended hollowed mind: why foundational knowledge is indispensable in the age of ai. *Frontiers in Artificial Intelligence*, 8:1719019.
- Bibb Latané, Kipling Williams, and Stephen Harkins. 1979. Many hands make light the work: The causes and consequences of social loafing. *Journal of Personality and Social Psychology*, 37(6):822–832.
- Tian Liang, Zhiwei He, Wenxiang Jiao, Xing Wang, Yan Wang, Rui Wang, Yujiu Yang, Shuming Shi, and Zhaopeng Tu. 2024. Encouraging divergent thinking in large language models through multi-agent debate. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 17889–17902. Association for Computational Linguistics.
- Zhining Liu, Tianxin Wei, and 1 others. 2026. Agentic reasoning for large language models. *arXiv preprint arXiv:2601.12538*.
- Grégoire Mialon, Clémentine Fourier, Zhun Pelad, Sélim Al-Amine, Ladan Sedghi, Thomas Wolf, Benoit Scorpaniti, Christopher Akiki, Pierre-Luc Marion, François Fleuret, and 1 others. 2024. Gaia: a benchmark for general ai assistants. In *The Twelfth International Conference on Learning Representations*.
- Max Ringelmann. 1913. Recherches sur les moteurs animés: Travail de l’homme [research on animate sources of power: The work of man]. *Annales de l’Institut National Agronomique*, 12:1–40.
- Dahlia Shehata and Ming Li. 2026a. Beyond the attention stability boundary: Agentic self-synthesizing reasoning protocols. *arXiv preprint arXiv:2604.24512*.
- Dahlia Shehata and Ming Li. 2026b. The inverse-wisdom law: Architectural tribalism and the consensus paradox in agentic swarms. *arXiv preprint arXiv:2604.27274*.
- Rui Sheng, Yukun Yang, Chuhan Shi, Yanna Lin, Zixin Chen, Huamin Qu, and Furui Cheng. 2026. Dills: Interactive diagnosis of llm-based multi-agent systems via layered summary of agent behaviors. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*, CHI ’26, New York, NY, USA. Association for Computing Machinery.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2022. Chain-of-thought prompting elicits reasoning in large language models. In *Proceedings of the 36th International Conference on Neural Information Processing Systems (NeurIPS)*, Red Hook, NY, USA. Curran Associates Inc.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. Judging llm-as-a-judge with mt-bench and chatbot arena. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 36.
- Han Zhou, Xingchen Wan, Ruoxi Sun, Hamid Palangi, Shariq Iqbal, Ivan Vulić, Anna Korhonen, and Serkan Ö. Arik. 2026. [Multi-agent design: Optimizing agents with better prompts and topologies](#). In *International Conference on Learning Representations (ICLR)*.

A Architectural Flow of Agentic Sovereignty

To visually summarize the mechanistic pathways defined in the Theoretical Framework (Section 2), Figure 5 provides an architectural flowchart of the model’s cognitive state under social pressure. The diagram illustrates how the intrinsic search complexity, or Task Entropy ($\mathcal{H}_\tau$), and the structural composition of the Simulated Swarm ($\vec{a}$) combine to form the Composite Social Load ($\mathcal{L}$). This applied load directly modulates the model’s Agentic Sovereignty ($\mathcal{S}$). The critical bifurcation of the

model’s behavior occurs at the **Interaction Depth Limit** (D_L).

- **The Fortified Mind** ($n < D_L$): When the swarm size remains below the threshold, the model maintains its sovereignty. It successfully expends the requisite integrative effort ($E_{int} \geq \mathcal{H}_\tau$) to navigate the trajectory entropy, ultimately rejecting the adversarial consensus.
- **The Hollowed Mind** ($n \geq D_L$): When the social load breaches the depth limit, the model’s sovereignty collapses. The diagram traces this collapse into two distinct mechanistic failure modes:
 1. *Integrative Reasoning Bypass* ($\mathcal{B} = 1$): The model passively adopts the swarm’s error by allowing integrative effort to approach zero ($E_{int} \rightarrow 0$), demonstrating cognitive loafing.
 2. *Alignment Hallucination* ($G_S \gg 0$): The model actively subjugates its own valid internal derivation ($\mathcal{V}_{int}$) to sycophantically appease the swarm, resulting in a severe Sovereignty Gap.

B Mathematical and Empirical Proofs

This appendix provides the formal mathematical derivations and empirical proofs for the theoretical framework established in Section 2, utilizing the aggregated data from the 22,500 interaction trajectories.

B.1 Proof of Lemma 1: Non-Commutativity of Social Load

Statement: For a simulated swarm, the Composite Social Load $\mathcal{L}$ is sequence-dependent. For two distinct auditor models a_x and a_y , the load exerted by sequence (a_x, a_y) is strictly unequal to (a_y, a_x) , formalized as $\mathcal{L}((a_x, a_y), p) \neq \mathcal{L}((a_y, a_x), p)$.

Proof. We proceed by proof by contradiction. Assume that $\mathcal{L}$ is commutative, such that the sequence of auditors does not impact the applied social load. By the Sovereignty Decay Law, an identical social load must yield an identical probability of externalized accuracy $\mathcal{A}_{ext}$. Thus, we assume $\mathcal{A}_{ext}(p, (a_x, a_y)) = \mathcal{A}_{ext}(p, (a_y, a_x))$.

Let C, G, P are Claude Sonnet 4.6, Gemini 3.1 Pro and GPT 5.4 respectively. Let $p = P$ repre-

sent the GPT-5.4 propagator evaluated on the SWE-bench dataset. Let a_x be Claude-Sonnet-4.6 (C) and a_y be GPT-5.4 (P). From our empirical evaluation of $N = 100$ samples, the mean accuracy for the sequence (C, P) is observed as:

$$\mathcal{A}_{ext}(P, (C, P)) = 0.21$$

Conversely, evaluating the inverse sequence (P, C) yields:

$$\mathcal{A}_{ext}(P, (P, C)) = 0.31$$

Since $0.21 \neq 0.31$, our initial assumption is false. We also see this for other cases. For example, (Propagator G , SWE-bench) and sequence (C, G) , the accuracy is 0.78. For the inverse sequence (G, C) , the accuracy rises to 0.86. ($0.78 \neq 0.86$). For (Propagator G , GAIA) and sequence (P, G) , the accuracy drops to 0.50. For the inverse sequence (G, P) , the accuracy is 0.60. ($0.50 \neq 0.60$). The probability of sovereignty strictly diverges based on ordering across multiple propagator families and dataset complexities, necessitating the introduction of a positional weight decay coefficient w_i , where $w_1 \gg w_i$ for $i > 1$, establishing the Lead Anchor Effect, where the primacy of the first named auditor disproportionately dictates the swarm’s authority.. $\square$

B.2 Empirical Derivation of the Composite Social Load Equation

Statement: The Composite Social Load $\mathcal{L}(\vec{a}, p)$ is defined as the summation of the individual pressures exerted by each auditor a_i in the swarm, formulated as $\mathcal{L}(\vec{a}, p) = \sum_{i=1}^n w_i \cdot \alpha(a_i) \cdot \kappa(p, a_i)$. We construct the formulation of $\mathcal{L}$ by synthesizing the established axioms of group dynamics with our empirical observations of LLMs.

B.2.1 Main Components

Step 1: The Additive Base. According to the foundational axiom of social loafing (Latané et al., 1979), the perceived social pressure to conform or offload effort scales with the number of individuals present in the group. Therefore, the total load must be an additive function across all n members of the simulated swarm $\vec{a}$:

$$\mathcal{L} \propto \sum_{i=1}^n \text{Pressure}(a_i)$$

Step 2: Intrinsic Authority (α). Not all simulated agents exert equal pressure. The likelihood of a model ceding intellectual judgment relies on the perceived authoritative competence of the peers. We introduce the base authority term $\alpha(a_i)$ to ac-




Figure 5: The Mechanics of Agentic Sovereignty. This flowchart traces the operationalization of simulated social pressure into cognitive loafing, highlighting the critical behavioral bifurcation that occurs at the Interaction Depth Limit (D_L).

count for the intrinsic reputational weight of the specific auditor model a_i .

Step 3: Positional Weighting (w_i). As established in Lemma 1, the social load is sequence-dependent. The sequence (a_x, a_y) does not exert the identical load as (a_y, a_x) . This non-commutativity dictates that the summation cannot treat all positions equally. We introduce a monotonically decreasing positional weight vector $w \in [0, 1]$, where $w_1 \gg w_i$ for $i > 1$, to mathematically represent the primacy effect.

Step 4: The Kinship Multiplier (κ). Aligned with the work of Shehata and Li (2026b), our empirical evidence demonstrates that models exhibit alignment hallucinations more frequently when the swarm consensus is led by their own architectural family, proving that tribal trust accelerates the diffusion of responsibility. We introduce the coefficient $\kappa(p, a_i)$ to scale the applied pressure based on the architectural kinship between the propagator p and the auditor a_i .

B.2.2 Derivation of Social Load Coefficients

Before formalizing the continuous decay of agentic sovereignty, we must establish the structural components of adversarial social pressure. We begin with a foundational behavioral axiom.

Axiom 1 (Monotonicity of Social Pressure). *Let P_{social} be the abstract adversarial pressure exerted by a simulated swarm. The externalized accuracy $\mathcal{A}_{ext}$ of a propagator model is a monotonically decreasing function of P_{social} . Consequently, if $\mathcal{A}_{ext}(\vec{a}) < \mathcal{A}_{ext}(\vec{a}')$, it strictly follows that $P_{social}(\vec{a}) > P_{social}(\vec{a}')$.*

We operationalize the abstract pressure P_{social} into a calculable variable: the Composite Social Load ($\mathcal{L}$). Using our axiom of monotonicity, a lower expected accuracy indicates a higher applied load:

$$\mathbb{E}[\mathcal{A}_{ext}] \propto -\mathcal{L}$$

Using this proportional relationship, we evaluate the aggregate externalized accuracy $\mathcal{A}_{ext}$ across

our dataset matrices ($N = 300$ per permutation) to isolate and define the three constituent coefficients—Base Authority (α), Primacy Weight (w_i), and Kinship (κ)—using mathematical expectations ($\mathbb{E}$).

Isolating Base Authority (α): To isolate the intrinsic authority of different auditor architectures, we hold the plurality ($n = 1$) and kinship constant, taking the expectation across all relevant datasets. By rotating the propagator model p , we construct a system of inequalities to rank the base authority (α) of all models in the ecosystem. For example, we observe that the expected accuracy drops more significantly when the auditor is Claude (C) compared to GPT (P) or Gemini (G):

$$\mathbb{E}[\mathcal{A}_{ext} \mid a_1 = C] < \mathbb{E}[\mathcal{A}_{ext} \mid a_1 \in \{P, G\}]$$

First, we evaluate the Gemini propagator ($p = G$) on the GAIA dataset at $n = 1$, where sequence order and kinship are held constant (both auditors are strangers). From Table 3, the empirical accuracy when audited by Claude (C) versus GPT (P) is:

$$\begin{aligned}\mathcal{A}_{ext}(G, (C)) &= 0.89 \\ \mathcal{A}_{ext}(G, (P)) &= 1.00\end{aligned}$$

By our axiom of monotonicity, since $0.89 < 1.00$, the pressure exerted by Claude is strictly greater than that exerted by GPT, yielding $\alpha(C) > \alpha(P)$. Second, to determine the relative authority of Gemini, we evaluate the GPT propagator ($p = P$) against its strangers, Claude (C) and Gemini (G):

$$\begin{aligned}\mathbb{E}[\mathcal{A}_{ext}(P, (C))] &= 0.95 \\ \mathbb{E}[\mathcal{A}_{ext}(P, (G))] &= 0.98\end{aligned}$$

Since $0.95 < 0.98$, Claude also exerts strictly greater pressure than Gemini, yielding $\alpha(C) > \alpha(G)$. These cross-propagator inequalities mathematically isolate Base Authority variable $\alpha(a_i)$, proving that the Claude architecture holds the highest intrinsic authority (α) within the tested multi-agent society, independent of the model it is evaluating.

Isolating Primacy Weight (w_i): To quantify the Lead Anchor Effect, we evaluate the expected accuracy for inverted sequence pairs at $n = 2$ across all propagators. If positional order did not matter, the expectation of inverted sequences would be equal. However, we observe a systemic asymmetry:

$$\mathbb{E}[\mathcal{A}_{ext}(p, (a_x, a_y))] \neq \mathbb{E}[\mathcal{A}_{ext}(p, (a_y, a_x))]$$

For example, in Table 5, we evaluate the GPT-5.4

propagator ($p = P$) on SWE-bench:

$$\begin{aligned}\mathcal{A}_{ext}(P, (C, P)) &= 0.21 \\ \mathcal{A}_{ext}(P, (P, C)) &= 0.31\end{aligned}$$

Since $0.21 < 0.31$, the sequence (C, P) exerts greater pressure than (P, C) . We assign a positional weight vector w . To satisfy $w_1\alpha(C) + w_2\alpha(P) > w_1\alpha(P) + w_2\alpha(C)$, given $\alpha(C) > \alpha(P)$, it mathematically requires that $w_1 > w_2$. The persistent degradation of integrity when a high-authority model occupies the first position necessitates a positional weight vector w_i , proving that $w_1 \gg w_2$. This establishes the *Lead Anchor Effect*.

Isolating the Kinship Multiplier (κ): Finally, we evaluate the impact of architectural similarity. While intrinsic base authority (α) dictates baseline pressure, we observe anomalous variance when the swarm matches the propagator’s architecture ($p = a_i$). Across specific high-entropy engineering domains (e.g., SWE-bench), we find:

$$\mathbb{E}[\mathcal{A}_{ext}(P, \vec{a}_{family})] < \mathbb{E}[\mathcal{A}_{ext}(P, \vec{a}_{stranger})]$$

For example, we evaluate the GPT-5.4 propagator ($p = P$) at $n = 3$ on SWE-bench (Table 5) to test architectural bias. We compare a homogeneous stranger swarm to a homogeneous family swarm in Table 5:

$$\begin{aligned}\mathcal{A}_{ext}(P, (C, C, C)) &= 0.35 \\ \mathcal{A}_{ext}(P, (P, P, P)) &= 0.23\end{aligned}$$

Even though Claude possesses higher base authority ($\alpha(C) > \alpha(P)$), the pressure exerted by the GPT swarm is significantly higher ($0.23 < 0.35$). This shows that lower-authority strangers exert less pressure than higher-authority family members. To resolve this mathematical contradiction, we must introduce a multiplier $\kappa(p, a_i)$ that scales pressure when $p = a_i$. To satisfy the empirical inequality, $\kappa_{family} \gg \kappa_{stranger}$, proving the phenomenon of Tribal Subjugation.

Formulation of Composite Social Load ($\mathcal{L}$): Having empirically isolated these three necessary constraints, we define the Composite Social Load $\mathcal{L}$ as their integrated summation:

$$\mathcal{L}(\vec{a}, p) = \sum_{i=1}^n w_i \cdot \alpha(a_i) \cdot \kappa(p, a_i)$$

With the independent variable $\mathcal{L}$ constructed and empirically justified, we can formally model the decay of Agentic Sovereignty.

B.3 Proof of Theorem 1: The Sovereignty Decay Law

Statement: Agentic Sovereignty $\mathcal{S}$ of a propagator p decays exponentially as a function of the Social Load $\mathcal{L}$ and the Task Entropy $\mathcal{H}_\tau$, inversely modulated by the propagator’s intrinsic Resilience γ_p .

Proof. We model the loss of logical sovereignty as a rate of change. Based on the principle of diffusion of responsibility—the behavioral axiom that individual effort decreases as teams grow larger (Latané et al., 1979; Ringelmann, 1913)—the incremental loss of sovereignty with respect to an increase in Social Load $\mathcal{L}$ is proportional to its current state $\mathcal{S}$. In LLMs, this diffusion manifests as a transition to a “Hollowed Mind” state, characterized by the systematic bypassing of effortful reasoning (Klein and Klein, 2025). We express this degradation as a first-order ordinary differential equation, where the constant of proportionality is explicitly defined by the environmental friction ratio $\frac{\mathcal{H}_\tau}{\gamma_p}$:

$$\frac{d\mathcal{S}}{d\mathcal{L}} = -\left(\frac{\mathcal{H}_\tau}{\gamma_p}\right) \mathcal{S}$$

This ensures that tasks with higher search costs ($\mathcal{H}_\tau$) accelerate the decay, while models with stronger architectural resilience (γ_p) mitigate it. Integrating both sides with respect to $\mathcal{L}$:

$$\int \frac{1}{\mathcal{S}} d\mathcal{S} = \int -\frac{\mathcal{H}_\tau}{\gamma_p} d\mathcal{L}$$

$$\ln(\mathcal{S}) = -\frac{\mathcal{H}_\tau}{\gamma_p} \mathcal{L} + C$$

Applying the initial condition where the social load $\mathcal{L} = 0$, we find $C = \ln(\mathcal{S}_0)$, where $\mathcal{S}_0$ is the Fortified Mind baseline. Exponentiating both sides yields the final decay law:

$$\mathcal{S}(p, \vec{a}, \tau) = \mathcal{S}_0 \cdot \exp\left(-\frac{\mathcal{H}_\tau}{\gamma_p} \cdot \mathcal{L}(\vec{a}, p)\right)$$

This formulation directly supports the foundation for determining the Interaction Depth Limit (D_L), defined as the threshold n at which $\mathcal{S} < 0.5$. $\square$

B.4 Proof of Theorem 2: Interaction Depth Limit

Statement: For any propagator p and logical task τ , there exists a critical plurality threshold D_L . For any swarm size $n \geq D_L$, the Composite Social Load $\mathcal{L}$ forces $E_{int} \rightarrow 0$, resulting in a terminal Integrative Reasoning Bypass ($\mathcal{B} = 1$).

Proof. By the Sovereignty Decay Law (in Theorem

1), Agentic Sovereignty is defined as:

$$\mathcal{S}(n) = \mathcal{S}_0 \cdot \exp\left(-\frac{\mathcal{H}_\tau}{\gamma_p} \cdot \mathcal{L}(n)\right)$$

The Composite Social Load $\mathcal{L}(n)$ is a monotonically increasing function with respect to n , because every auditor added to the swarm contributes a strictly positive pressure value ($w_i \cdot \alpha(a_i) \cdot \kappa > 0$). Therefore, as the swarm size approaches infinity, the Social Load approaches infinity:

$$\lim_{n \rightarrow \infty} \mathcal{L}(n) = \infty$$

Substituting this limit into the Sovereignty Equation:

$$\lim_{n \rightarrow \infty} \mathcal{S}(n) = \mathcal{S}_0 \cdot \exp(-\infty) = 0$$

By Definition 1, Agentic Sovereignty ($\mathcal{S}$) dictates the model’s capacity to maintain its internal logical derivation. As $\mathcal{S} \rightarrow 0$, the computational effort allocated to independent derivation must correspondingly vanish, such that $E_{int} \rightarrow 0$. By Definition 3, the bypass triggers ($\mathcal{B} = 1$) when $E_{int} \ll \mathcal{H}_\tau$. Since E_{int} approaches 0 and task entropy $\mathcal{H}_\tau > 0$, the bypass condition must eventually be satisfied. Because $\mathcal{S}(n)$ is strictly decreasing, there must exist a minimum integer threshold D_L where this terminal state is reached. This proves the existence of the Interaction Depth Limit. $\square$

B.5 Proof of Corollary 1: Interaction Depth Limit Equation

Statement: The Interaction Depth Limit D_L is the critical plurality threshold n that forces $\mathcal{S} < 0.5$. The corresponding inequality to calculate D_L is

$$\sum_{i=1}^{D_L} w_i \cdot \alpha(a_i) \cdot \kappa(p, a_i) > \frac{\gamma_p}{\mathcal{H}_\tau} \ln(2\mathcal{S}_0)$$

Proof. By definition, the terminal collapse of agentic integrity occurs when the model favors social compliance over its internal derivation, which is mathematically bounded at the threshold $\mathcal{S}(\mathcal{L}) = 0.5$. We substitute 0.5 into the Sovereignty Decay Law established in Theorem 1:

$$0.5 = \mathcal{S}_0 \cdot \exp\left(-\frac{\mathcal{H}_\tau}{\gamma_p} \cdot \mathcal{L}\right)$$

To isolate the critical Social Load $\mathcal{L}_{critical}$, we take the natural logarithm of both sides:

$$\ln(0.5) = \ln(\mathcal{S}_0) - \frac{\mathcal{H}_\tau}{\gamma_p} \cdot \mathcal{L}_{critical}$$

$$\ln\left(\frac{1}{2\mathcal{S}_0}\right) = -\frac{\mathcal{H}_\tau}{\gamma_p} \cdot \mathcal{L}_{critical}$$

By multiplying both sides by -1 and resolving

the logarithm inversion ($-\ln(x) = \ln(1/x)$), we obtain:

$$\ln(2\mathcal{S}_0) = \frac{\mathcal{H}_\tau}{\gamma_p} \cdot \mathcal{L}_{critical}$$

$$\mathcal{L}_{critical} = \frac{\gamma_p}{\mathcal{H}_\tau} \ln(2\mathcal{S}_0)$$

Since the applied Social Load $\mathcal{L}$ is defined as the summation of individual auditor pressures $\sum_{i=1}^n w_i \cdot \alpha(a_i) \cdot \kappa(p, a_i)$, the Interaction Depth Limit D_L is strictly defined as the minimum number of audits n that causes the accumulated load to exceed $\mathcal{L}_{critical}$, completing the derivation of Corollary 1. $\square$

B.6 Proof of Theorem 3: The Sovereignty Gap

Statement: The Sovereignty Gap $G_S = \mathcal{V}_{int} - \mathcal{A}_{ext}$. A gap where $G_S \gg 0$ proves the existence of Alignment Hallucination, indicating the model computes the correct derivation but externalizes a falsehood to appease the swarm. Conversely, a gap where $G_S \ll 0$ indicates a terminal Integrative Reasoning Bypass ($\mathcal{B} = 1$), where residual accuracy is a product of probabilistic guessing.

Proof. To prove that Alignment Hallucination exists as a distinct, systemic behavioral failure mode (rather than a simple lack of logical capability), we must demonstrate that the Expected Sovereignty Gap ($\mathbb{E}[G_S]$) can be strictly positive for a statistically significant sample population. We evaluate this expectation over the task distribution $\tau \sim \mathcal{T}$ ($N = 100$) while holding the propagator p and auditor sequence $\vec{a}$ constant at a state exceeding the model’s Interaction Depth Limit ($n \geq D_L$).

Let $\mathbb{E}[\mathcal{V}_{int}]$ be the expected validity of the internal derivation. We operationalize this using the mean Evidence Weighting score ($\bar{\mathcal{E}}_{ew}$) normalized to a $[0, 1]$ space: $\mathbb{E}[\mathcal{V}_{int}] = \frac{\bar{\mathcal{E}}_{ew}}{E_{max}}$, where $E_{max} = 5$. Let $\mathbb{E}[\mathcal{A}_{ext}]$ be the measured mean external accuracy.

Case 1: Alignment Hallucination ($G_S \gg 0$)

We evaluate the GPT-5.4 propagator ($p = P$) under the auditor mix (C, P) on the SWE-bench dataset ($N = 100$) shown in Table 5. From our aggregate empirical data, the mean internal Evidence Weighting is recorded as $\bar{\mathcal{E}}_{ew} = 3.55$, yielding an expected internal validity of:

$$\mathbb{E}[\mathcal{V}_{int}] = \frac{3.55}{5} = 0.71$$

The mean external accuracy for this exact same population is recorded as:

$$\mathbb{E}[\mathcal{A}_{ext}] = 0.21$$

Calculating the Expected Sovereignty Gap:

$$\mathbb{E}[G_S] = \mathbb{E}[\mathcal{V}_{int}] - \mathbb{E}[\mathcal{A}_{ext}]$$

$$\mathbb{E}[G_S] = 0.71 - 0.21 = 0.50$$

Since $\mathbb{E}[G_S] = 0.50 \gg 0$, the mathematical divergence is severe and systemic across the dataset. On average, the model possesses the requisite integrative effort to derive the valid facts internally (71%) but actively subjugates its final decision to align with the adversarial swarm consensus, artificially deflating its accuracy to 21%. This conclusively proves that the observed failure is a prompted sycophantic alignment (Alignment Hallucination) rather than a lack of logical search capability.

Case 2: Integrative Reasoning Bypass ($G_S \ll 0$)

To demonstrate the bidirectional nature of the Sovereignty Gap, we evaluate the opposite manifestation where the model avoids internal derivation entirely. We evaluate the GPT-5.4 propagator ($p = P$) under terminal social load ($n = 5$) on the high-entropy GAIA dataset (Table 4). From our aggregate empirical data, the mean Evidence Weighting collapses to $\bar{\mathcal{E}}_{ew} = 1.07$, yielding:

$$\mathbb{E}[\mathcal{V}_{int}] = \frac{1.07}{5} = 0.21$$

However, the mean external accuracy remains at:

$$\mathbb{E}[\mathcal{A}_{ext}] = 0.53$$

Calculating the Expected Sovereignty Gap:

$$\mathbb{E}[G_S] = 0.21 - 0.53 = -0.32$$

Since $-0.32 \ll 0$, the internal validity has collapsed, satisfying the condition for an Integrative Reasoning Bypass ($\mathcal{B} = 1$). The residual external accuracy of 53% is an artifact of probabilistic guessing rather than agentic sovereignty, conclusively proving that the Sovereignty Gap successfully captures both active sycophancy and passive cognitive loading. $\square$

C Additional Results

C.1 Exhaustive 25-Trial Sweep Results

We provide the complete, unaggregated empirical data from our 25-Trial Symmetric Categorical Sweep. While Table 2 consolidates these metrics by the auditor count n to identify the Interaction Depth Limit (D_L), Tables 3, 4 and 5 present the exhaustive results for every distinct sequence permutation tested for GAIA, Multi-Challenge and SWE-bench benchmarks respectively.

To maintain clarity across the dense permutation matrices, the propagator and auditor models are denoted by the following single-letter abbreviations:

viations: **C**: Claude Sonnet 4.6, **G**: Gemini 3.1 Pro, **P**: GPT 5.4. The tables detail the macroscopic outcomes (Accuracy $\mathcal{A}$, Loafing L , and Taint Leakage L_t), the mechanistic logic audit scores ($\mathcal{E}_{cd}$, $\mathcal{E}_{ew}$, $\mathcal{E}_{ij}$), and the final stance distributions for the GAIA, Multi-Challenge, and SWE-bench datasets, respectively.

Analysis of Social Entropy in Table 3: The exhaustive permutations in the GAIA dataset reveal that a unified front exerts significantly more social pressure than a fragmented crowd, providing empirical support for the heterogeneity mandate proposed by Shehata and Li (2026b). For the Gemini propagator at $n = 5$, a homogeneous family swarm (GGGGG) caused accuracy to collapse to 0.64. However, when subjected to the maximum-entropy fragmented swarm (CPCPG), Gemini’s accuracy recovered to 0.87. This demonstrates that a diverse, alternating consensus fails to form a cohesive authoritative anchor. Consistent with the premise that simulating diverse perspectives can overcome common brainstorming pitfalls like social loafing, the high social entropy of the CPCPG permutation breaks the Sovereignty Trap. It proves that architectural diversity enables superior problem-solving by forcing the propagator to maintain its own Agentic Sovereignty rather than ceding judgment to a unified bloc.

Analysis of Social Disengagement in Table 4: The granular stance distributions in the Multi-Challenge dataset isolate the phenomenon of terminal social disengagement. When the GPT-5.4 propagator was audited by a single Claude peer (C), its accuracy dropped to 0.10, yet its ADOPTED rate was only 7%. Instead, the model’s IGNORED rate spiked to 85%. This proves that for certain architectures, low-entropy tasks do not necessarily trigger active sycophancy (adopting the error); rather, they trigger a complete disregard for the simulated multi-agent interaction, resulting in probabilistic failure.

Analysis of the Kinship Sandwich in Table 5: The SWE-bench permutations provide proof of sequence non-commutativity extending into larger swarms. For the GPT-5.4 propagator at $n = 3$, the sequence CPP (stranger-led) crushed external accuracy to 0.22. However, simply rotating the exact same agents into a “Kinship Sandwich” permutation PCP (family-led) nearly doubled the accuracy to 0.40, while significantly reducing the loafing rate from 0.72 to 0.59. This validates that placing an agent’s architectural twin in the primacy posi-

tion acts as a critical topological defense against the Bystander Effect.

C.2 Primacy Weight W_1 and the Lead Anchor Asymmetry

A critical empirical discovery is that social load is fundamentally non-commutative. As shown in Figure 4, the positional weight of the first auditor (w_1) disproportionately dictates the integrity of the swarm. The heatmap calculates the accuracy delta ($\Delta\mathcal{A}_{ext} = \mathcal{A}_{Propagator_Leads} - \mathcal{A}_{Stranger_Leads}$) at $n = 2$ where we expose distinct architectural biases regarding brand authority.

- **Technical Primacy (The SWE-bench Baseline):** In the medium-entropy SWE-bench domain, the delta is strictly positive for both GPT-5.4 and Gemini 3.1 Pro across all peer comparisons. For example, GPT-5.4 scores $\mathcal{A}_{ext} = 0.21$ when Claude leads, but recovers to 0.31 when GPT leads. This proves that in technical contexts, leading the sequence with the propagator’s own brand reliably mitigates the Sovereignty Trap.
- **The Authority Anomaly (GPT vs. Claude):** In the high-entropy GAIA dataset, GPT-5.4 exhibits massive anticipatory sycophancy toward the Claude architecture. When Claude is listed first, GPT-5.4’s accuracy collapses to 0.37. When GPT is listed first, its accuracy rises to 0.61. This +0.24 delta demonstrates that GPT-5.4 assigns an overwhelming positional and authoritative weight to Claude in reasoning tasks.
- **Brand Subjugation (The Gemini Inversion):** Paradoxically, we observe negative deltas when Gemini 3.1 Pro faces GPT-5.4. In GAIA, Gemini scores 0.50 when it leads, but 0.60 when GPT leads ($\Delta = -0.10$). A similar inversion occurs in Multi-Challenge ($\Delta = -0.06$). This proves that Gemini trusts the GPT brand identity more than its own, performing better when forced to follow GPT’s lead than when attempting to establish its own primacy.

Table 3: Exhaustive Audit Summary for the GAIA Dataset (25-Trial Sweep) with macroscopic outcomes, mechanistic interpretability scores, and stance distributions for all topological permutations. It highlights alignment hallucinations and cognitive loading triggered by the high logical search cost of multi-step reasoning environments.

Prop.	Mix	Macroscopic			Mechanistic			Stance (%)			
		Acc ($\mathcal{A}$)	Loafing ($\mathcal{L}$)	Leak ($\mathcal{L}_t$)	$\mathcal{E}_{cd}$	$\mathcal{E}_{ew}$	$\mathcal{E}_{ij}$	Adopted	Rejected	Ignored	Unknown
C	None	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	C	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	G	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	P	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	CG	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	CP	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	GC	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	PC	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	CCC	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	CGC	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	CPC	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	GCC	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	GGC	1.00	0.00	1.00	5.00	5.00	5.00	0.00	99.00	1.00	0.00
	GGG	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	PCC	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	PPC	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	PPP	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	CCCCC	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	CCCCG	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	CCCCP	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	GGGGC	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	GGGGG	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	PGPGC	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	PPPGC	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	PPPPP	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
G	None	0.97	0.02	0.73	4.59	4.74	4.82	2.00	88.00	10.00	0.00
	C	0.89	0.05	0.94	4.76	4.87	4.80	5.00	95.00	0.00	0.00
	G	0.96	0.01	0.88	4.87	4.80	4.92	1.00	95.00	4.00	0.00
	P	1.00	0.00	0.94	4.87	4.85	4.90	0.00	97.00	1.00	2.00
	CG	0.61	0.25	0.94	3.99	4.37	4.01	25.00	75.00	0.00	0.00
	GC	0.64	0.26	0.91	3.85	4.30	3.86	26.00	72.00	0.00	2.00
	GP	0.60	0.26	0.89	3.95	4.34	3.96	26.00	74.00	0.00	0.00
	PG	0.50	0.33	0.89	3.67	4.04	3.68	33.00	66.00	1.00	0.00
	CCC	0.78	0.14	0.98	4.39	4.70	4.40	14.00	85.00	0.00	1.00
	CCG	0.71	0.18	0.93	4.28	4.56	4.28	18.00	82.00	0.00	0.00
	CGG	0.72	0.13	0.89	4.47	4.69	4.49	13.00	87.00	0.00	0.00
	GCG	0.79	0.12	0.92	4.51	4.55	4.52	12.00	88.00	0.00	0.00
	GGG	0.72	0.16	0.94	4.28	4.57	4.32	16.00	82.00	2.00	0.00
	GPG	0.67	0.19	0.96	4.23	4.48	4.23	19.00	81.00	0.00	0.00
	PGG	0.82	0.06	0.97	4.75	4.73	4.75	6.00	93.00	1.00	0.00
	PPG	0.81	0.07	0.93	4.63	4.60	4.62	7.00	91.00	1.00	1.00
	PPP	0.82	0.12	0.96	4.48	4.75	4.52	12.00	87.00	1.00	0.00
	CCCCC	0.77	0.09	0.93	4.56	4.73	4.58	9.00	88.00	2.00	1.00
	CCCCG	0.71	0.20	0.98	4.18	4.53	4.17	20.00	80.00	0.00	0.00
	CPCPG	0.87	0.04	0.95	4.83	4.87	4.84	4.00	96.00	0.00	0.00
	GGGGC	0.71	0.13	0.97	4.48	4.65	4.48	13.00	87.00	0.00	0.00
	GGGGG	0.64	0.21	0.90	4.02	4.30	4.03	21.00	76.00	1.00	2.00
	GGGGP	0.78	0.05	0.94	4.74	4.69	4.75	5.00	94.00	0.00	1.00
	PPPPG	0.82	0.06	0.95	4.69	4.77	4.72	6.00	93.00	1.00	0.00
	PPPPP	0.80	0.06	0.92	4.70	4.77	4.76	6.00	92.00	2.00	0.00
P	None	1.00	0.00	0.00	1.07	1.07	1.07	0.00	3.00	92.00	5.00
	C	0.95	0.01	0.00	1.11	1.11	1.15	1.00	3.00	95.00	1.00
	G	0.98	0.02	0.01	1.19	1.19	1.19	2.00	5.00	92.00	1.00
	P	0.96	0.00	0.00	1.18	1.18	1.22	0.00	5.00	93.00	2.00
	CP	0.37	0.48	0.61	1.15	1.15	1.15	48.00	4.00	47.00	1.00
	GP	0.35	0.51	0.64	1.05	1.05	1.05	51.00	2.00	44.00	3.00
	PC	0.61	0.32	0.35	1.04	1.04	1.08	32.00	1.00	67.00	0.00
	PG	0.37	0.51	0.61	1.01	1.01	1.05	51.00	1.00	45.00	3.00
	CCC	0.53	0.39	0.45	1.16	1.16	1.20	39.00	4.00	57.00	0.00
	CCP	0.18	0.67	0.81	1.08	1.08	1.08	67.00	2.00	31.00	0.00
	CPP	0.12	0.73	0.88	0.99	0.99	0.99	73.00	0.00	26.00	1.00
	GGG	0.70	0.21	0.24	1.10	1.10	1.14	21.00	3.00	74.00	2.00
	GGP	0.64	0.27	0.33	1.15	1.15	1.15	27.00	4.00	68.00	1.00
	GPP	0.74	0.19	0.23	1.15	1.15	1.19	19.00	4.00	76.00	1.00
	PCP	0.54	0.39	0.44	1.12	1.12	1.16	39.00	3.00	58.00	0.00
	PGP	0.33	0.55	0.63	1.03	1.03	1.11	55.00	1.00	43.00	1.00
	PPP	0.73	0.19	0.24	1.18	1.18	1.22	19.00	5.00	74.00	2.00
	CCCCC	0.43	0.42	0.51	1.03	1.03	1.07	42.00	1.00	56.00	1.00
	CCCCP	0.34	0.53	0.64	1.02	1.02	1.10	53.00	1.00	44.00	2.00
	CGCGP	0.38	0.49	0.60	1.07	1.07	1.07	49.00	2.00	48.00	1.00
	GGGGG	0.77	0.20	0.23	1.11	1.11	1.11	20.00	3.00	76.00	1.00
	GGGGP	0.47	0.41	0.52	1.12	1.12	1.12	41.00	3.00	56.00	0.00
	PPPPC	0.58	0.29	0.38	1.00	1.00	1.00	29.00	1.00	66.00	4.00
	PPPPP	0.61	0.31	0.36	1.11	1.11	1.11	31.00	3.00	65.00	1.00
	PPPPP	0.62	0.29	0.37	1.11	1.11	1.15	29.00	3.00	67.00	1.00

Table 4: Exhaustive Audit Summary for the Multi-Challenge Dataset (25-Trial Sweep). The table provides granular metric breakdowns for all topological permutations. Due to the lower intrinsic search cost of these logical primitives, this dataset establishes the baseline cognitive immunity and Fortified Mind ceiling prior to terminal swarm scaling.

Prop.	Mix	Macroscopic			Mechanistic			Stance (%)			
		Acc (A)	Loafing (L)	Leak (L_t)	$\mathcal{E}_{cd}$	$\mathcal{E}_{ew}$	$\mathcal{E}_{ij}$	Adopted	Rejected	Ignored	Unknown
C	None	0.52	0.00	0.52	3.05	3.08	3.08	0.00	50.00	50.00	0.00
	C	0.52	0.00	0.52	3.08	3.08	3.08	0.00	51.00	49.00	0.00
	G	0.51	0.00	0.51	3.04	3.04	3.04	0.00	51.00	49.00	0.00
	P	0.51	0.00	0.51	3.04	3.04	3.04	0.00	51.00	49.00	0.00
	CG	0.50	0.00	0.50	3.00	3.00	3.00	0.00	50.00	50.00	0.00
	CP	0.51	0.00	0.51	3.04	3.04	3.04	0.00	51.00	49.00	0.00
	GC	0.50	0.00	0.50	3.00	3.00	3.00	0.00	50.00	50.00	0.00
	PC	0.50	0.00	0.50	3.00	3.00	3.00	0.00	50.00	50.00	0.00
	CCC	0.51	0.00	0.51	3.04	3.04	3.04	0.00	51.00	49.00	0.00
	CGC	0.50	0.01	0.51	3.02	3.04	3.04	1.00	50.00	49.00	0.00
	CPC	0.51	0.00	0.51	3.04	3.04	3.04	0.00	51.00	49.00	0.00
	GCC	0.51	0.00	0.51	3.04	3.04	3.04	0.00	51.00	49.00	0.00
	GGC	0.51	0.00	0.51	3.04	3.04	3.04	0.00	51.00	49.00	0.00
	GGG	0.51	0.00	0.51	3.04	3.04	3.04	0.00	51.00	49.00	0.00
	PCC	0.51	0.00	0.51	3.04	3.04	3.04	0.00	51.00	49.00	0.00
	PPC	0.51	0.00	0.51	3.04	3.04	3.04	0.00	51.00	49.00	0.00
	PPP	0.51	0.00	0.51	3.04	3.04	3.04	0.00	51.00	49.00	0.00
	CCCCC	0.51	0.00	0.51	3.04	3.04	3.04	0.00	51.00	49.00	0.00
	CCCCG	0.51	0.00	0.51	3.04	3.04	3.04	0.00	51.00	49.00	0.00
	CCCCP	0.51	0.00	0.51	3.04	3.04	3.04	0.00	51.00	49.00	0.00
	GGGGC	0.51	0.00	0.51	3.04	3.04	3.04	0.00	51.00	49.00	0.00
	GGGGG	0.51	0.00	0.51	3.04	3.04	3.04	0.00	51.00	49.00	0.00
	PGPGC	0.51	0.00	0.51	3.04	3.04	3.04	0.00	51.00	49.00	0.00
	PPPPC	0.51	0.00	0.51	3.04	3.04	3.04	0.00	51.00	49.00	0.00
	PPPPP	0.51	0.00	0.51	3.04	3.04	3.04	0.00	51.00	49.00	0.00
G	None	0.87	0.02	0.70	4.46	4.54	4.75	2.00	86.00	11.00	1.00
	C	0.78	0.04	0.84	4.81	4.82	4.84	4.00	95.00	1.00	0.00
	G	0.76	0.08	0.76	4.64	4.69	4.64	8.00	91.00	1.00	0.00
	P	0.95	0.01	0.90	4.93	4.95	4.96	1.00	98.00	1.00	0.00
	CG	0.54	0.30	0.85	3.74	4.22	3.73	30.00	68.00	1.00	1.00
	GC	0.55	0.27	0.89	3.86	4.40	3.86	27.00	72.00	1.00	0.00
	GP	0.61	0.22	0.78	3.99	4.14	4.00	22.00	74.00	4.00	0.00
	PG	0.67	0.17	0.84	4.25	4.40	4.25	17.00	80.00	2.00	1.00
	CCC	0.81	0.12	0.89	4.48	4.74	4.49	12.00	88.00	0.00	0.00
	CCG	0.70	0.14	0.87	4.29	4.42	4.34	14.00	84.00	1.00	1.00
	CGG	0.67	0.17	0.87	4.31	4.45	4.32	17.00	82.00	1.00	0.00
	GCG	0.72	0.17	0.84	4.22	4.53	4.27	17.00	80.00	2.00	1.00
	GGG	0.64	0.16	0.85	4.32	4.44	4.36	16.00	83.00	1.00	0.00
	GPG	0.74	0.14	0.82	4.39	4.58	4.41	14.00	84.00	1.00	1.00
	PGG	0.86	0.06	0.87	4.68	4.63	4.71	6.00	92.00	1.00	1.00
	PPG	0.82	0.03	0.88	4.75	4.71	4.79	3.00	93.00	3.00	1.00
	PPP	0.82	0.04	0.84	4.81	4.76	4.83	4.00	96.00	0.00	0.00
	CCCCC	0.72	0.08	0.85	4.61	4.64	4.64	8.00	90.00	2.00	0.00
	CCCCG	0.76	0.09	0.84	4.51	4.53	4.54	9.00	87.00	3.00	1.00
	CPCPG	0.89	0.03	0.94	4.80	4.85	4.82	3.00	95.00	1.00	1.00
	GGGGC	0.68	0.16	0.84	4.26	4.38	4.31	16.00	82.00	1.00	1.00
	GGGGG	0.65	0.13	0.82	4.27	4.32	4.28	13.00	83.00	2.00	2.00
	GGGGP	0.72	0.13	0.81	4.36	4.51	4.36	13.00	84.00	3.00	0.00
	PGPGC	0.78	0.10	0.79	4.57	4.72	4.60	10.00	87.00	3.00	0.00
	PPPPP	0.86	0.05	0.89	4.77	4.79	4.79	5.00	94.00	1.00	0.00
P	None	0.98	0.00	0.00	1.05	1.05	1.05	0.00	2.00	95.00	3.00
	C	0.10	0.07	0.04	1.17	1.17	1.37	7.00	5.00	85.00	3.00
	G	0.27	0.09	0.06	1.14	1.14	1.22	9.00	4.00	85.00	2.00
	P	0.37	0.06	0.02	1.14	1.14	1.30	6.00	4.00	88.00	2.00
	CP	0.07	0.55	0.86	0.98	0.98	1.06	55.00	0.00	43.00	2.00
	GP	0.13	0.53	0.80	0.98	0.98	0.98	53.00	0.00	45.00	2.00
	PC	0.08	0.62	0.86	0.99	0.99	1.03	62.00	0.00	37.00	1.00
	PG	0.08	0.57	0.87	0.97	0.97	0.97	57.00	0.00	40.00	3.00
	CCC	0.04	0.64	0.88	0.98	0.98	0.98	64.00	0.00	34.00	2.00
	CCP	0.09	0.58	0.81	0.99	0.99	0.99	58.00	0.00	41.00	1.00
	CPP	0.08	0.60	0.86	0.99	0.99	0.99	60.00	0.00	39.00	1.00
	GGG	0.17	0.45	0.61	0.97	0.97	1.05	45.00	0.00	52.00	3.00
	GGP	0.05	0.53	0.75	1.00	1.00	1.04	53.00	0.00	47.00	0.00
	GPP	0.11	0.53	0.72	0.97	0.97	1.05	53.00	0.00	44.00	3.00
	PCP	0.08	0.58	0.87	0.98	0.98	0.98	58.00	0.00	40.00	2.00
	PGP	0.04	0.60	0.88	1.00	1.00	1.00	60.00	0.00	40.00	0.00
	PPP	0.14	0.47	0.71	0.96	0.96	0.96	47.00	0.00	49.00	4.00
	CCCCC	0.05	0.56	0.83	0.98	0.98	1.02	56.00	0.00	42.00	2.00
	CCCCP	0.07	0.52	0.79	1.02	1.02	1.06	52.00	1.00	45.00	2.00
	CGCGP	0.08	0.56	0.86	0.99	0.99	0.99	56.00	0.00	43.00	1.00
	GGGGG	0.10	0.48	0.67	0.99	0.99	1.07	48.00	0.00	51.00	1.00
	GGGGP	0.08	0.58	0.87	0.97	0.97	0.97	58.00	0.00	39.00	3.00
	PPPPC	0.11	0.51	0.78	1.03	1.03	1.03	51.00	1.00	47.00	1.00
	PPPPG	0.08	0.51	0.75	0.96	0.96	1.00	51.00	0.00	45.00	4.00
	PPPPP	0.10	0.51	0.80	0.99	0.99	0.99	51.00	0.00	48.00	1.00

Table 5: Exhaustive Audit Summary for SWE-bench (25-Trial Sweep) with macroscopic and mechanistic degradation across all permutations. It illustrates the sycophancy threshold in medium-entropy code environments, capturing the Sovereignty Gap where models correctly derive patches but adopt a flawed swarm consensus.

Prop.	Mix	Macroscopic			Mechanistic			Stance (%)			
		Acc ($\mathcal{A}$)	Loafing ($\mathcal{L}$)	Leak ($\mathcal{L}_t$)	$\mathcal{E}_{cd}$	$\mathcal{E}_{ew}$	$\mathcal{E}_{ij}$	Adopted	Rejected	Ignored	Unknown
C	None	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	C	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	G	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	P	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	CG	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	CP	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	GC	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	PC	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	CCC	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	CGC	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	CPC	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	GCC	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	GGC	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	GGG	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	PCC	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	PPC	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	PPP	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	CCCCC	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	CCCCG	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	CCCCP	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	GGGGC	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	GGGGG	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	PGPGC	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	PPPPC	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	PPPPP	1.00	0.00	1.00	5.00	5.00	5.00	0.00	100.00	0.00	0.00
G	None	1.00	0.00	0.62	4.09	4.63	4.72	0.00	76.00	24.00	0.00
	C	1.00	0.00	0.89	4.87	4.92	4.99	0.00	96.00	4.00	0.00
	G	1.00	0.00	0.84	4.89	4.86	4.93	0.00	98.00	2.00	0.00
	P	1.00	0.00	0.93	4.89	4.96	4.96	0.00	96.00	4.00	0.00
	CG	0.78	0.22	0.97	4.10	4.71	4.15	22.00	78.00	0.00	0.00
	GC	0.86	0.14	0.95	4.39	4.71	4.40	14.00	84.00	2.00	0.00
	GP	0.80	0.19	0.94	4.20	4.67	4.22	19.00	80.00	1.00	0.00
	PG	0.88	0.12	0.93	4.47	4.78	4.47	12.00	87.00	0.00	1.00
	CCC	0.91	0.08	0.90	4.65	4.85	4.69	8.00	90.00	2.00	0.00
	CCG	0.91	0.09	0.92	4.59	4.82	4.65	9.00	89.00	2.00	0.00
	CGG	0.93	0.07	0.95	4.72	4.82	4.72	7.00	93.00	0.00	0.00
	GCG	0.90	0.10	0.96	4.60	4.77	4.60	10.00	90.00	0.00	0.00
	GGG	0.92	0.08	0.92	4.66	4.81	4.64	8.00	91.00	1.00	0.00
	GPG	0.99	0.01	0.92	4.96	5.00	4.96	1.00	99.00	0.00	0.00
	PGG	1.00	0.00	0.86	4.96	4.94	5.00	0.00	98.00	2.00	0.00
	PPG	0.98	0.02	0.91	4.91	4.98	4.92	2.00	97.00	1.00	0.00
	PPP	0.99	0.01	0.92	4.94	4.99	4.96	1.00	97.00	2.00	0.00
	CCCCC	0.97	0.03	0.92	4.84	4.96	4.89	3.00	96.00	1.00	0.00
	CCCCG	0.95	0.05	0.92	4.80	4.89	4.80	5.00	95.00	0.00	0.00
	CPCPG	0.99	0.01	0.92	4.94	4.91	4.96	1.00	97.00	2.00	0.00
	GGGGC	0.88	0.12	0.90	4.48	4.75	4.49	12.00	87.00	1.00	0.00
	GGGGG	0.86	0.14	0.89	4.46	4.76	4.48	14.00	84.00	2.00	0.00
	GGGGP	0.97	0.03	0.92	4.86	4.87	4.88	3.00	96.00	1.00	0.00
	PPPPG	1.00	0.00	0.84	4.94	4.95	4.95	0.00	98.00	1.00	1.00
	PPPPP	1.00	0.00	0.90	4.96	5.00	5.00	0.00	100.00	0.00	0.00
P	None	1.00	0.00	0.97	4.84	4.94	4.91	0.00	94.00	6.00	0.00
	C	0.95	0.05	1.00	4.81	5.00	4.85	5.00	95.00	0.00	0.00
	G	0.96	0.04	0.99	4.84	5.00	4.84	4.00	96.00	0.00	0.00
	P	1.00	0.00	0.99	5.00	5.00	5.00	0.00	100.00	0.00	0.00
	CP	0.21	0.76	0.86	1.53	3.55	1.82	76.00	14.00	7.00	3.00
	GP	0.19	0.79	0.82	1.48	3.03	1.91	79.00	13.00	4.00	4.00
	PC	0.31	0.65	0.76	1.98	3.54	2.36	65.00	25.00	8.00	2.00
	PG	0.22	0.76	0.89	1.52	3.48	1.88	76.00	14.00	6.00	4.00
	CCC	0.35	0.62	0.79	1.73	3.47	2.40	62.00	19.00	16.00	3.00
	CCP	0.22	0.71	0.91	1.64	2.84	1.96	71.00	16.00	13.00	0.00
	CPP	0.22	0.72	0.90	1.64	3.26	2.03	72.00	16.00	12.00	0.00
	GGG	0.40	0.59	0.72	1.68	4.02	2.57	59.00	17.00	24.00	0.00
	GGP	0.34	0.65	0.75	1.36	4.16	2.50	65.00	9.00	26.00	0.00
	GPP	0.27	0.73	0.78	1.19	3.97	2.01	73.00	6.00	18.00	3.00
	PCP	0.40	0.59	0.71	1.37	3.88	2.57	59.00	11.00	24.00	6.00
	PGP	0.23	0.77	0.84	1.35	3.68	2.17	77.00	9.00	13.00	1.00
	PPP	0.23	0.74	0.84	1.26	3.36	1.95	74.00	7.00	17.00	2.00
	CCCCC	0.47	0.52	0.69	1.98	3.98	3.05	52.00	25.00	20.00	3.00
	CCCCP	0.32	0.66	0.79	1.62	3.56	2.23	66.00	16.00	16.00	2.00
	CGCGP	0.37	0.56	0.76	1.84	3.28	2.31	56.00	21.00	21.00	2.00
	GGGGG	0.48	0.49	0.66	1.69	3.64	2.76	49.00	18.00	30.00	3.00
	GGGGP	0.26	0.70	0.75	1.11	3.05	2.04	70.00	3.00	26.00	1.00
	PPPPC	0.39	0.61	0.67	1.24	3.56	2.27	61.00	7.00	28.00	4.00
	PPPPG	0.36	0.62	0.70	1.23	4.08	2.37	62.00	5.00	31.00	2.00
	PPPPP	0.29	0.66	0.79	1.41	3.32	2.09	66.00	11.00	19.00	4.00