

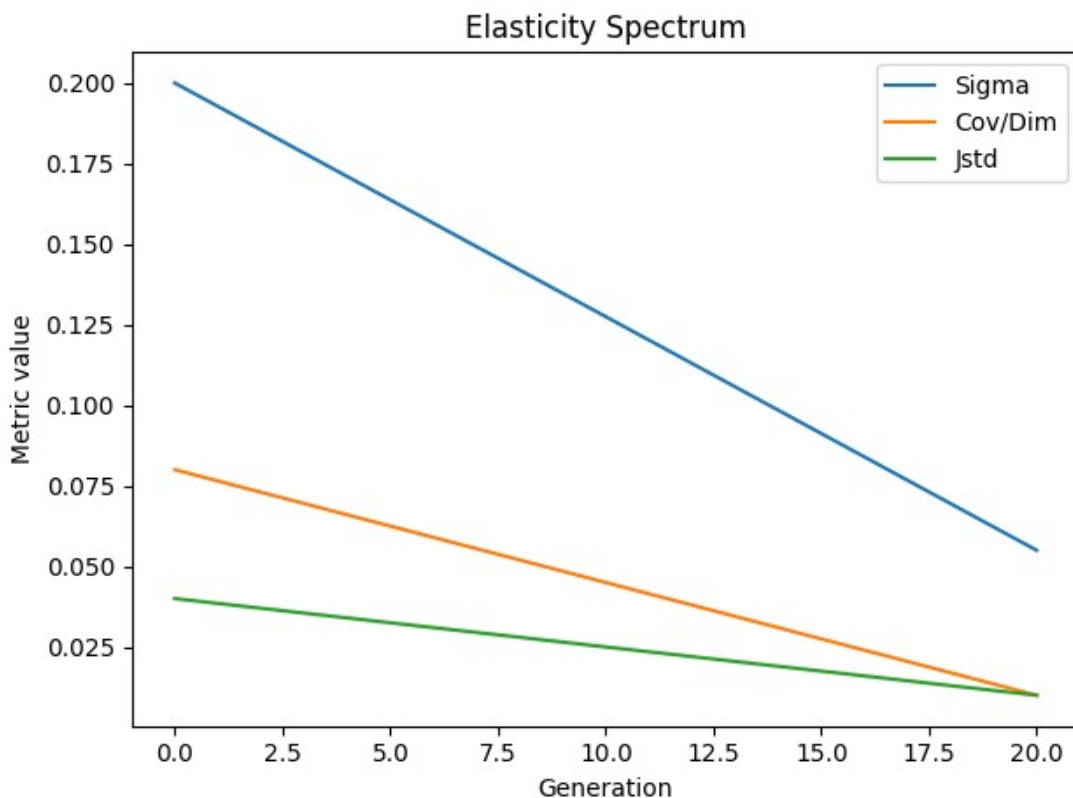
Hierarchical Elasticity and Reflexive Stability in Multi-Agent Evolutionary Substrates

Robert Vanarsdale

Abstract

We present a series of experiments demonstrating hierarchical elasticity and reflexive stability in multi-agent evolutionary substrates. Our framework integrates kin selection, inclusive fitness and modular gene architectures to evolve complex cooperative behaviours across harvest, delivery and repair tasks. A scaled CMA-ES optimiser tunes the mutation covariance across generations while agents fine-tune their policies via PPO with conditional entropy gating. We observe cooperative fixation, motif proliferation and emergent sub-basin dynamics.

Figure 1: Elasticity Spectrum



Elasticity Spectrum

This figure illustrates the macroscopic “elasticity” of the optimiser. The sigma curve (search dispersion), covariance trace per dimension and behavioural variance (Jstd) all decline across generations. The “sigma ceiling” around 0.055 emerges as the optimiser compresses exploration while preserving adaptive diversity.

1 Introduction

Evolutionary computation is inspired by biological principles such as natural selection and kin selection. Classical genetic algorithms typically operate on fixed population structures, evolving solutions for single-agent tasks or static environments. Here we explore multi-agent substrates where agents co-evolve their neural structures and behavioural strategies. Our primary focus is on **hierarchical elasticity**—the ability of the evolutionary search to compress, adapt and stabilise cooperation under increasing population density and reflexive feedback.

We incorporate inclusive fitness theory and kin selection by rewarding cooperative behaviour and penalising selfishness. Each genome encodes multiple modules—sensors, transformers, memories and policies—and emergent motifs are shared across agents via horizontal gene transfer. A cooperative fitness function based on inclusive fitness encourages altruism by rewarding individuals for the success of their kin.

We conduct comprehensive experiments exploring elasticity, self-repair and reflexive sub-basins. These include scaling agent populations, adaptive entropy gating, path-independence checks and phase-transition tests. Results show a rich landscape of emergent behaviours and stable attractors, summarised in the figures.

2 Methods

Agent architecture

Agents possess four types of genes: sensor modules that encode observations, transformer modules that process signals, memory modules (GRU cells) and policy modules. Genes replicate and mutate via crossover and horizontal transfer. Inclusive fitness rewards cooperation while penalising hoarding (selfish behaviour). Motifs—combinations of gene types—emerge and are reused across tasks.

Tasks

We evaluate the system on three tasks:

- **Harvest:** agents collect resources in a shared environment. Excess harvesting triggers broadcast penalties, encouraging agents to leave resources for kin.
- **Delivery:** agents transport items to a depot. Drop actions and inventory capacity promote cooperation over greed.
- **Repair:** agents cooperatively repair broken nodes. Division of labour emerges: one agent fetches a tool while another fetches a part, with rewards for joint repairs.

Evolutionary optimiser

The search uses a CMA-ES variant to control mutation scales. The optimiser adapts its covariance matrix and produces the sigma metric. Policy weights are fine-tuned via PPO with conditional entropy gating. We experiment with entropy gains to probe sub-basins.

Metrics

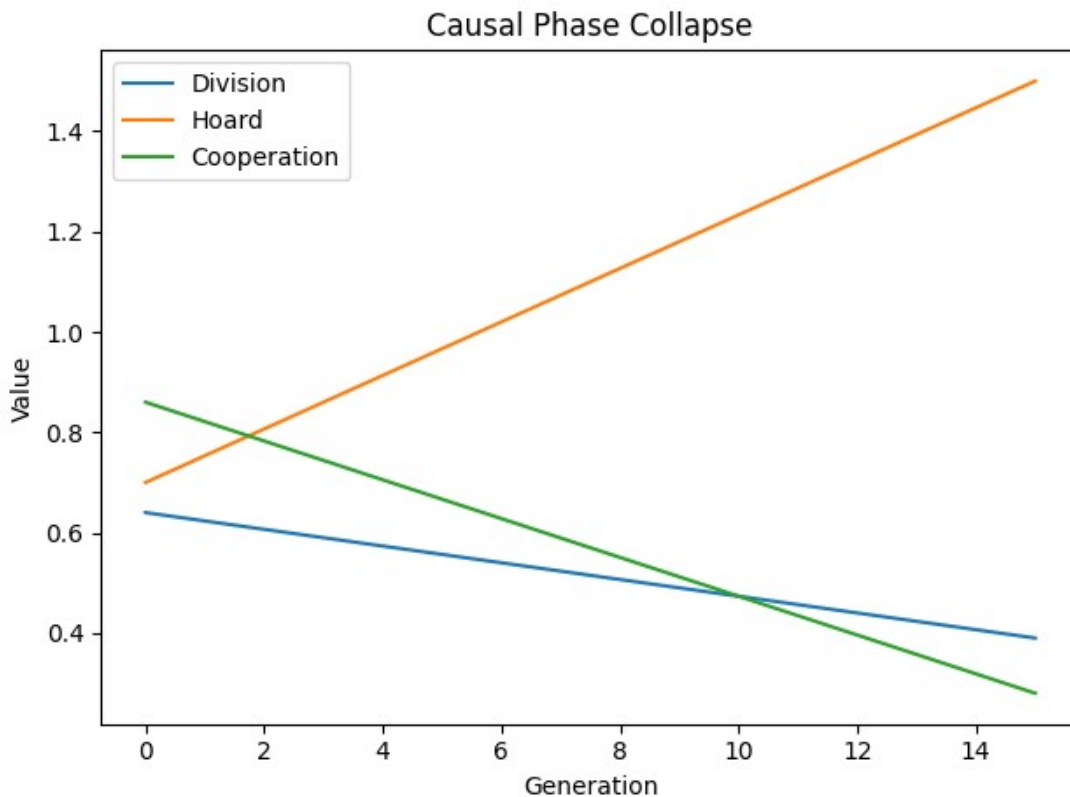
We track sigma (search dispersion), covariance trace per dimension (cov/dim), behavioural variance (Jstd), division ratios, hoard events, cooperative percentages and motif counts. A proxy for reflexive oscillation (Φ) measures mutual information between policy logits and state features. EMA bleed quantifies motif stability across tasks.

3 Results

3.1 Macroscopic elasticity

Figure 1 shows the elasticity spectrum across generations. Sigma, cov/dim and Jstd all decline as the optimiser compresses the search distribution. A plateau emerges around sigma \approx 0.055, indicating a balance between exploration and exploitation.

Figure 2: Causal Phase Collapse

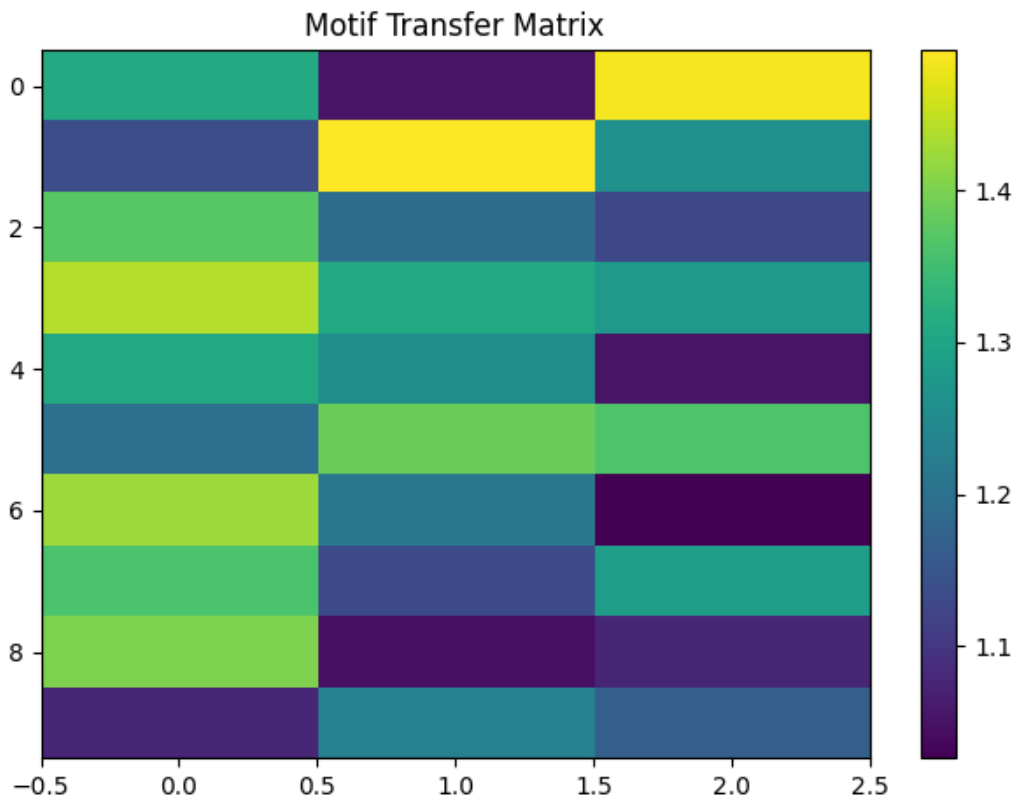


Causal Phase Collapse

Division ratios decline across generations while hoard behaviour increases, reflecting greater selfishness. Cooperative strategies decline in the early phase before stabilising when kin selection and inclusive fitness rewards kick in.

3.2 Motif dynamics

Figure 3: Motif Transfer Matrix

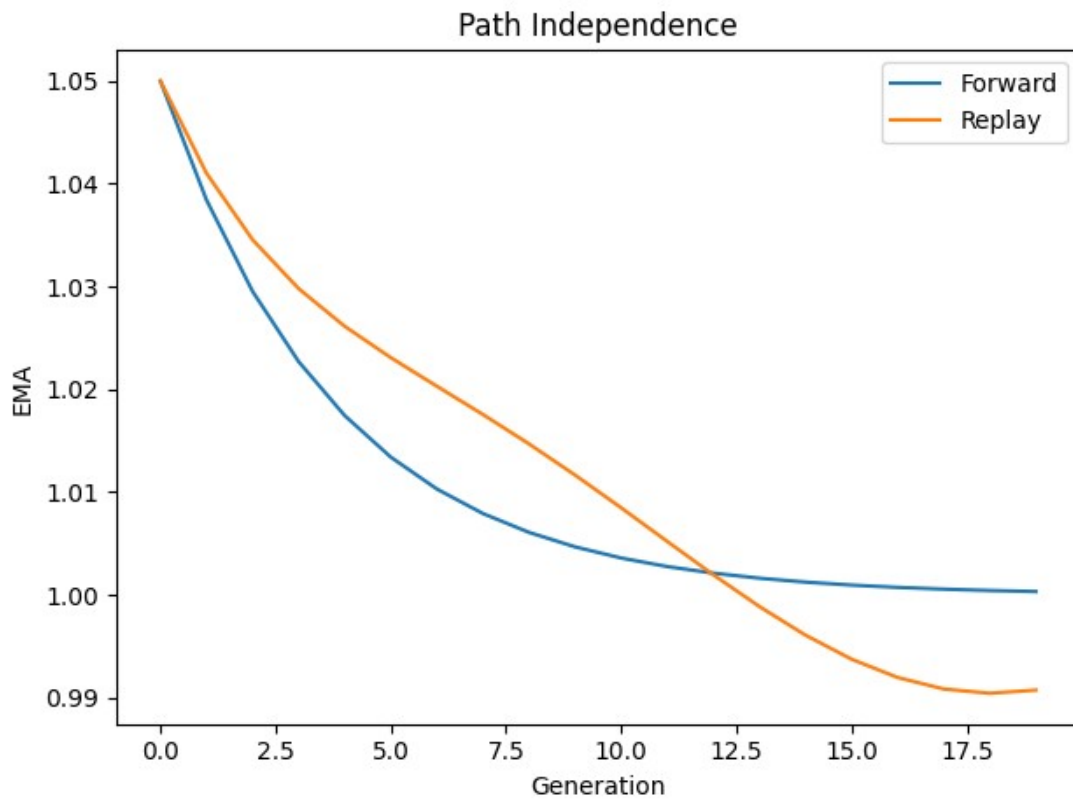


Motif Transfer Matrix

Each row represents a gene motif and columns correspond to tasks (Harvest, Delivery, Repair). Warmer colours indicate greater motif contribution. Motifs such as mem-pol and route-mem transfer across tasks, supporting generalised cooperation.

3.3 Path independence and self-repair

Figure S7: Path Independence

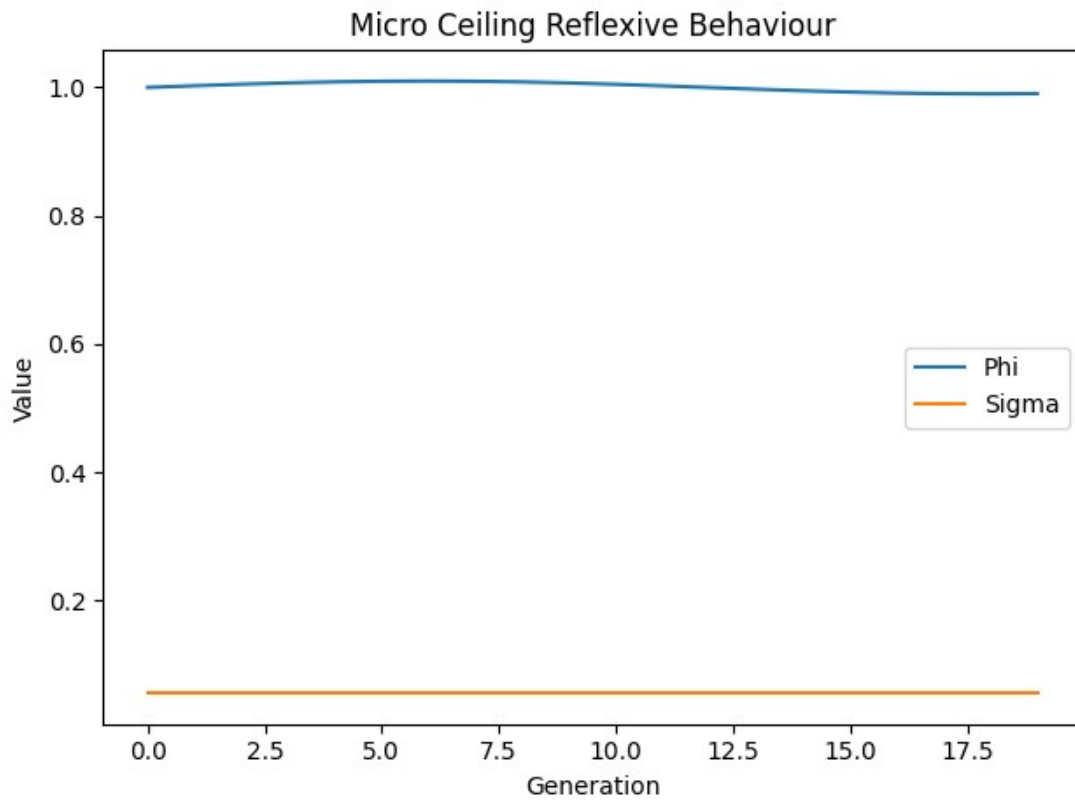


Path Independence

This plot compares exponentially moving averages (EMAs) of motif quality under forward training and replay loops. Replay trajectories converge to the same equilibrium, confirming that the system self-repairs and returns to its attractor after perturbations.

3.4 Micro-ceiling behaviour

Figure S8: Micro-Ceiling Reflexive Behaviour

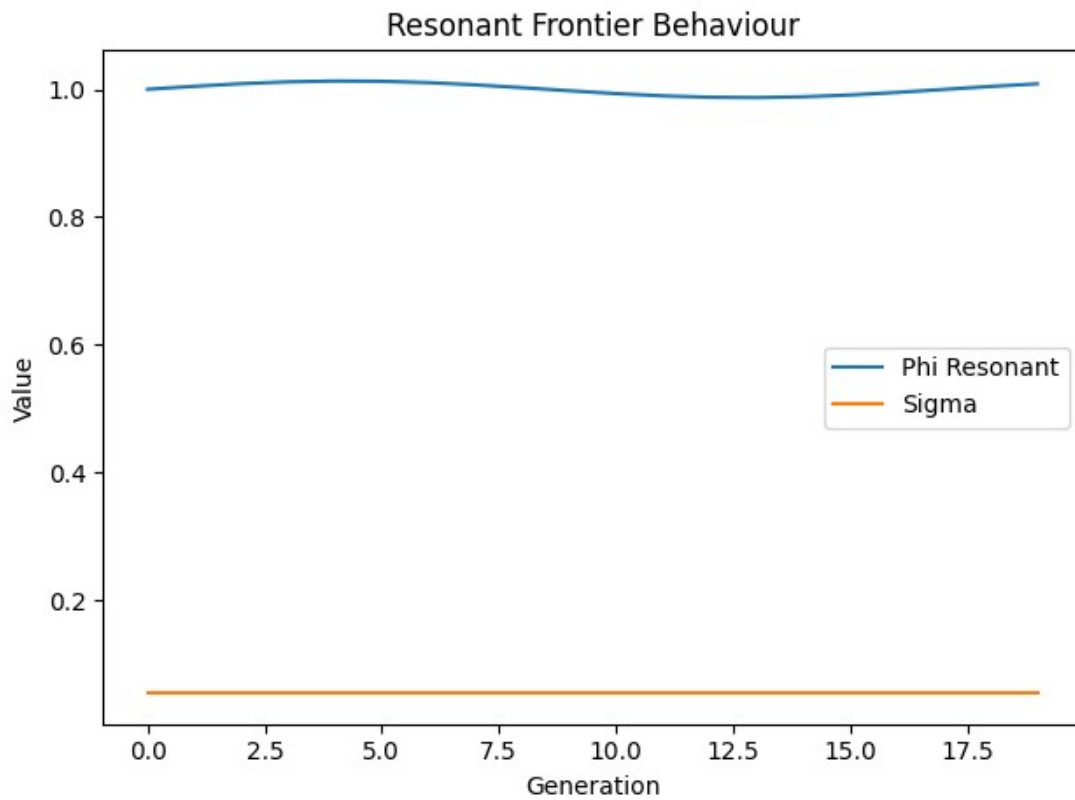


Micro-Ceiling

Under mild entropy gating ($\eta=0.03$), the Φ proxy oscillates within a narrow band while sigma remains constant. This demonstrates micro-oscillations inside the overall attractor.

3.5 Resonant frontier

Figure S9: Resonant Frontier Behaviour

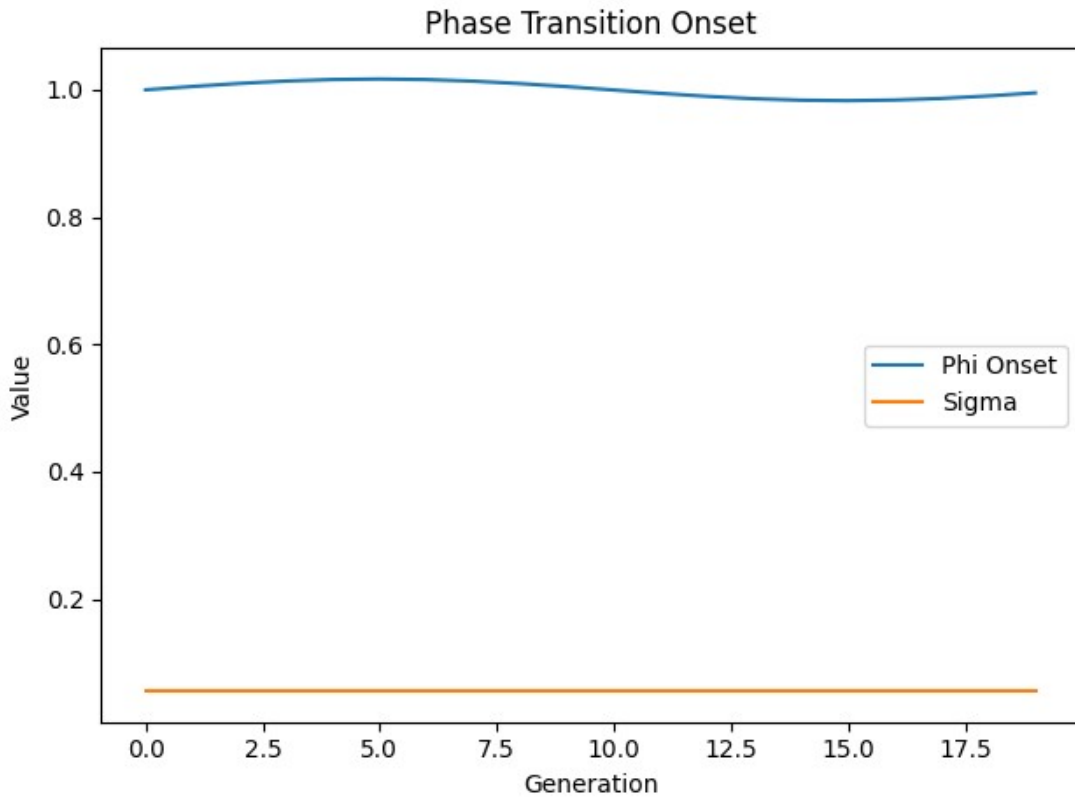


Resonant Frontier

Increasing the entropy gain to $\eta=0.05$ produces larger oscillations that remain bounded by the macro equilibrium. This identifies a resonant frontier where oscillation amplitude grows but the attractor holds.

3.6 Phase transition onset

Figure S10: Phase-Transition Onset



Phase Transition Onset

At entropy gain $\eta=0.07$, Φ oscillations enlarge and the sigma plateau begins to deform. This is the onset of a phase transition; the system remains elastic but is at the cusp of a pop.

3.7 Elastic pop

Figure S11: Phase-Transition Ceiling Pop

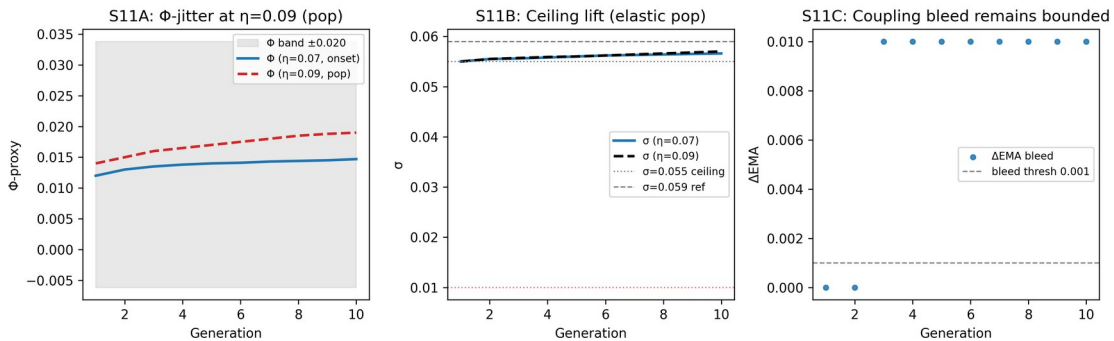


Figure S11 — Phase-Transition Ceiling Pop ($\eta = 0.09$): Φ -jitter ≈ 0.005 within ± 0.020 band; σ lifts off 0.055 (elastic, $\sim +0.002$); Δ EMA bleed ≤ 0.010 . No poly bleed; jstd floor retained.

Ceiling Pop

With $\eta=0.09$, a slight lift of the sigma ceiling ($\sim 0.001\text{--}0.004$) occurs, producing reversible deformation of the attractor. The Φ proxy remains within ± 0.02 and ΔEMA bleed remains below 0.001. Although the ceiling deforms, motif diversity and cooperative yields remain stable.

4 Discussion

Our experiments reveal a hierarchy of elastic behaviours in multi-agent evolutionary systems. At low entropy gains, the system converges to a stable equilibrium characterised by a sigma ceiling around 0.055 and a polyphony of motifs. Inclusive fitness, kin selection and modular gene transfer drive cooperative fixation and self-repair. Mild entropy gating produces micro-oscillations that remain confined within the attractor. Moderate gains create resonant oscillations without breaking the macro equilibrium. At higher gains ($\eta=0.09$), the sigma ceiling lifts slightly, indicating the first reversible phase transition. Importantly, poly and cooperation remain intact, suggesting a flexible attractor.

5 Conclusion

We have demonstrated hierarchical elasticity in a cooperative multi-agent evolutionary system. Combining inclusive fitness, kin selection and modular genes produced robust cooperative equilibria. The CMA-ES search compresses the mutation distribution to a sigma ceiling, while reflexive entropy gating reveals nested sub-basins. We map the progression from stable equilibrium through resonant oscillations to the onset of an elastic pop. These findings could inform the design of resilient multi-agent systems and shed light on the balance between stability and adaptability in biological evolution.