

From Cooperation to Hierarchy: A Study of Dynamics of Hierarchy Emergence in a Multi-Agent System

Shanshan Mao¹ and Peter Tino¹

School of Computer Science, University of Birmingham, Birmingham, UK
sxm1550@student.bham.ac.uk, p.tino@bham.ac.uk

Abstract. A central premise in evolutionary biology is that individual variation can generate information asymmetries that facilitate the emergence of hierarchical organisation. To examine this process, we develop an agent-based model (ABM) to identify the minimal conditions under which hierarchy arises in dynamic multi-agent systems, focusing on the roles of initial heterogeneity and mutation amplitude across generations. Hierarchical organisation is quantified using the Trophic Incoherence (TI) metric, which captures directional asymmetries in interaction networks. Our results show that even small individual differences can be amplified through repeated local interactions involving reproduction, competition, and cooperation, but that hierarchical order is markedly more sensitive to mutation amplitude than to initial heterogeneity. Across repeated trials, stable hierarchies reliably emerge only when mutation amplitude is sufficiently high, while initial heterogeneity primarily affects early formation rather than long-term persistence. Overall, these findings demonstrate how simple interaction rules can give rise to both the emergence and persistence of hierarchical organisation, providing a quantitative account of how structured inequality can develop from initially homogeneous populations.

Keywords: hierarchy emergence · trophic incoherence · multi-agent systems

1 Introduction and Background

From cooperative hunting in African social animals [6, 2] to the stratification of modern human societies [15], hierarchical structures—often implicit [9]—remain tightly linked to task distribution and resource allocation [28, 22]. Such recurring patterns of structured coordination, observed across natural and artificial systems [29], are commonly interpreted as manifestations of hierarchy. Yet how hierarchy arises from initially homogeneous populations, and whether small trait variations can trigger abrupt transitions toward structured inequality, remains an open question. Against this backdrop, we examine how decentralised coordination and reinforcement dynamics, together with a minimal set of ingredients, can generate stable hierarchical organisation in complex adaptive systems.

In this context, hierarchy is not understood as a fixed social ranking or institutional ordering, but as a structural property emerging from interaction patterns. More specifically, we focus on the formation of persistent directional asymmetries in information and influence flow among agents, arising endogenously through repeated local interactions.

From a modelling perspective, although many theoretical approaches examine how hierarchical structures can arise from initially equivalent individuals [1, 23, 27, 12], comparatively few provide quantitative frameworks that capture both the stochastic onset of differentiation and the subsequent consolidation of advantage within a unified evolutionary setting, with most existing approaches emphasising either ecological or social mechanisms in isolation.

An important exception is the work of Perret and colleagues, who demonstrate that interactions between individual heterogeneity and group size can shape the emergence of dominance hierarchies in patch-based environments [20, 19]. Their framework provides a useful conceptual basis for the present study.

Building on this line of work, we introduce an agent-based modelling (ABM) framework [26] for a biologically inspired, decentralised multi-agent system.

Within this framework, agents repeatedly interact to express internal capabilities and negotiate transient leadership roles [7, 24, 8]. The model integrates ecological, social, and evolutionary processes—combining density-dependent reproduction, stochastic mortality, speaker–listener coordination, and probabilistic resource allocation—to reproduce feedback loops underlying hierarchical organisation.

More broadly, the concept of hierarchy varies across disciplines. In complex adaptive systems, it often denotes nested, multi-level subsystems that organise information and control flows [15, 4], while socio-economic studies emphasise statistical regularities such as rank–size scaling [3, 5]. Despite these differences, both perspectives emphasise differentiation and asymmetry as central features of complex systems.

Consistent with the interactional perspective adopted here, hierarchy is therefore treated functionally as the *directionality of information flow* among agents. A system is more hierarchical when influence propagates predominantly from higher to lower levels, and less so when interactions are largely reciprocal. This interpretation emphasises asymmetric organisation rather than fixed positional ranking.

To quantify this interactional notion of hierarchy, we employ the *Trophic Incoherence* (TI) framework [13, 21, 25], a continuous measure of directional order in networks originally developed for ecological food webs. Low TI values correspond to more ordered, hierarchical structures, whereas higher values indicate flatter or more disordered interaction patterns. TI therefore provides a principled diagnostic for identifying the onset and persistence of hierarchy in decentralised populations.

Within this framework, we identify the minimal conditions under which hierarchical order can arise and stabilise by systematically varying initial heterogeneity and mutation amplitude. Adopting Trophic Incoherence as a structural measure further enables the temporal tracking of hierarchy emergence without presupposing fixed leadership or ranking mechanisms.

2 Methodology and Experimental Design of the Self-Organising Cooperation System

Our agent-based model (ABM) [26] examines the emergence of stable hierarchies from local interactions within a decentralised population, using a minimal environment in which simple behavioural rules give rise to structured, group-level outcomes. Each

agent maintains internal states—energy, spatial position, behavioral mode, and an α -value representing social standing—and may act independently or in coordination with others. Agents move, regulate energy, and seek partners; when cooperating, they perform collective foraging and reproduction [11, 17], subject to proximity, energy sufficiency, and behavioral compatibility. Although cooperation is common, collective activities also introduce competition: access to resources shapes reproductive success, while occupying a leadership role yields greater influence and returns. Together, these interactions generate selection pressures that couple survival and reproduction with both individual traits and social behaviour. As a consequence, agents with higher α -values more frequently assume the speaker role, which in turn reinforces their standing and forms the basis of our experimental setup for examining the emergence of hierarchical structure in a decentralised multi-agent system.

2.1 Self-Organisation Principles

Beyond individual behaviour, the system incorporates ecological feedbacks that regulate population-level stability, with reproduction and survival jointly responding to population density and resource availability [30, 10]. Population growth follows a standard density-dependent mechanism: as the number of agents N_t approaches the carrying capacity K , reproduction slows due to resource limitation and increases again at lower densities. The birth rate at time t is defined as

$$BR_t = f \cdot \left(1 - \frac{N_t}{K}\right), \quad (1)$$

where $f \in (0, 1)$ denotes the intrinsic fertility parameter, influenced by agent attributes such as age or oestrus status. This formulation imposes a smooth upper bound on growth, ensuring a gradual decline in reproduction near environmental capacity.

Mortality follows a Poisson-based process [30], with risk increasing for older or critically energy-depleted agents. Together, these feedbacks create a competitive yet stable ecological setting in which survival and reproduction depend on both internal agent states and overall population density. In the baseline configuration, death events occur stochastically, with higher probability for agents exceeding an age threshold ($\tau_t^i \geq 10$) or with zero energy ($E_t^i = 0$), independent of hierarchical position.

2.2 State Representation and Behavior Dynamics

In the Multi-Agent System (MAS), each agent interacts with both its environment and other agents through a state-driven process involving sensing, decision-making, and action execution.

Fig. 1 illustrates the overall architecture of agent–environment interactions. To formally represent these interactions, we define the state of the system at each discrete time t as a vector $\underline{S}_t$ that aggregates the states of all agents in the set $\mathcal{N}$:

$$\underline{S}_t = (S_t^1, S_t^2, \dots, S_t^N), \quad (2)$$




Fig. 1: Agent–environment interaction diagram adapted from [16, 18], illustrating the information flow between agent states, sensing, and action execution within the MAS.

where S_t^i represents the state of agent i at time t , and N is the total number of agents in the system. This formulation provides a complete mapping of the system configuration at each discrete time interval. Here, the superscript indicates the agent index, while the subscript denotes the time index throughout the subsequent equations.

The environment evolves in response to agent actions and system states. In our model, the environment includes spatial features such as resource locations (food), reproduction zones (villages), and agent proximity fields, all of which influence agents' decisions. Formally, the environmental state at time $t + 1$ is given by the environmental response function $T_E(\cdot)$:

$$\underline{\mathcal{E}}_{t+1} = T_E(\underline{S}_t, \underline{\mathcal{E}}_t, \underline{A}_t), \quad (3)$$

where $\underline{\mathcal{E}}_t$ denotes relevant environmental states at time t , and $\underline{A}_t$ denotes the actions performed by agents at time t . This function captures how collective actions and current conditions jointly shape environmental dynamics. Similarly, the state transition for each agent is determined by the current agent states, environmental conditions, and actions undertaken by agents. The agent state transition function $T_S(\cdot)$ is defined as:

$$\underline{S}_{t+1} = T_S(\underline{S}_t, \underline{\mathcal{E}}_t, \underline{A}_t). \quad (4)$$

The transition captures how agent–environment interactions update individual states over time. In the proposed Multi-Agent System (MAS), each agent determines its next action at discrete time step t through an agent action function $G(\cdot)$:

$$A_{t+1}^i = G(\underline{S}_t, \underline{\mathcal{E}}_t, \underline{A}_t), \quad (5)$$

where A_{t+1}^i denotes the action of agent i at time $t + 1$. The function $G(\cdot)$ maps current system states, environmental conditions, and past actions to future agent behaviors.

The nature of $G(\cdot)$ determines whether agent behavior is stochastic or deterministic. All agent actions are updated synchronously at each discrete time step. If $G(\cdot)$ incorporates probabilistic elements—such as stochastic decision rules or sampling from softmax distributions—the resulting actions A_{t+1}^i exhibit stochastic characteristics. Similarly, when the environmental update function $T_E(\cdot)$ includes random or uncertain components, both the environmental state $\underline{\mathcal{E}}_{t+1}$ and subsequent agent states $\underline{S}_{t+1}$ will also evolve stochastically.

When both the agent function $G(\cdot)$ and the environmental function $T_E(\cdot)$ operate deterministically, the system's evolution becomes fully predictable. This structure makes it possible to compare stochastic and deterministic variants of the model under the same update rules. Each agent at time t maintains a set of internal state variables, including energy level $E_t^i \in \mathbb{R}_{\geq 0}$, sensed environmental information, proximity to resources and settlements, and the availability of neighboring agents for cooperative behavior. These variables collectively guide the agent's decision-making process and give rise to a range of behavioral modes, which can be categorized into individual and joint actions.

- **Single-Agent Actions:** These actions do not require querying the states or previous actions of other agents. They include:
 - **DEAD:** Triggered if $E_t^i = 0$ or the agent's age $\tau_t^i \geq 10$, at which point the risk of death increases with age. The age of agent i is represented by $\tau_t^i \in \mathbb{N}$, measured in discrete time steps.
 - **MOVING_TO_VILLAGE:** Occurs when an agent has sufficient energy ($E_t^i > 10$) but is located beyond a threshold distance R_{village} from a village, which could be represented as $Dv_t^i \geq R_{\text{village}}^m$. Dv_t^i denotes the Euclidean distance from agent i to the boundary of village m , with $m \in \{1, 2, 3, 4, 5\}$ and R_{village}^m the corresponding village radius.
 - **IDLE:** Default, lowest-priority action when no other conditions are met.
 - **WAIT_REPRO:** if the agent's energy is more than 10 ($E_t^i > 10$) and its position is inside of one village $Dv_t^i \leq R_{\text{village}}^m$.
- **Joint Actions:** These actions involve explicit interactions with and queries of the states or previous actions of other agents:
 - **REPRODUCTION (2 agents):** Two agents collaborate to reproduce, contingent upon energy levels, proximity to a village, and synchronized previous action states. Agent's sex is encoded as a binary variable, $\text{sex}^i \in \{\text{male}, \text{female}\}$. Mathematically, reproduction occurs only if:

$$\text{REPRO}_{t+1}^{(i,j)} = \begin{cases} 1, & \begin{aligned} &\text{if } E_t^i > 10 \text{ and } E_t^j > 10 \\ &\text{and } Dv_t^i \leq R_{\text{village}}^m \\ &\text{and } Dv_t^j \leq R_{\text{village}}^m \\ &\text{and } A_t^i = A_t^j = \text{WAIT_REPRO} \\ &\text{and } \text{sex}^i \neq \text{sex}^j \end{aligned} \\ 0, & \text{otherwise.} \end{cases} \quad (6)$$

Reproductive success further depends on agents' relative α -scores, influencing offspring number and reinforcing hierarchical dynamics.

- **COOPERATION (3 agents):** Three agents collaboratively perform tasks like food collection tasks. Initiating such tasks requires agents to evaluate their energy level, nearby idle agents, and available food resources. The collective task initiation criterion is mathematically expressed as:

$$\text{COOP}_{t+1}^{(i,j,k)} = \begin{cases} 1, & \begin{aligned} &\text{if } E_t^i < 10 \text{ and } v_t^i \geq 2 \\ &\text{and } \mu_t^i \geq 1, \end{aligned} \\ 0, & \text{otherwise.} \end{cases} \quad (7)$$

where ν_t^i denotes the number of nearby IDLE agents, which $\nu_t^i \in \mathbb{N}$ typically does not exceed 5, and μ_t^i denotes the number of accessible food resources (within radius ρ of the agent).

2.3 Inheritance and Mutation in Reproduction

Following reproduction, a successful pairing between two agents produces one or more offspring whose capabilities arise from a biologically inspired inheritance–mutation process balancing trait transmission and stochastic variability.

At initialization, each agent’s innate capability C_i is drawn from $\mathcal{N}(100, \text{parameter}_c)$, where parameter_c controls population heterogeneity. Larger parameter_c values yield broader capability spreads, while smaller ones produce more uniform populations.

Let $C^{p1} > 0$ and $C^{p2} > 0$ denote the parental capacities of agents $p1$ and $p2$. The offspring’s inherited capability C^k is a weighted combination of the parental maximum and mean, biased toward the stronger parent to model selective inheritance:

$$C^k = h \cdot \max(C^{p1}, C^{p2}) + (1 - h) \cdot \frac{C^{p1} + C^{p2}}{2}, \quad (8)$$

where $h \in [0, 1]$ is the heritability coefficient, set to $h = 0.9$ to represent strong vertical transmission from the fitter parent with minor contribution from the other. Equation (8) is bounded by parental values, while stochastic mutation introduced in Eq. (9) prevents monotonic amplification across generations.

To introduce variability, a mutation mechanism perturbs the inherited value with probability p_m by a zero-mean Gaussian noise term:

$$C^{k'} = C^k + \delta, \quad \text{where } \delta \sim \mathcal{N}(0, \sigma^2), \quad (9)$$

where σ denotes the mutation scale (parameter_μ , referred to as the *mutation amplitude* in the Results). Increasing parameter_μ amplifies variance and evolutionary exploration, whereas smaller values maintain conservative inheritance. If no mutation occurs, $C^{k'} = C^k$.

Through this process, mutation continuously reintroduces heritable variation across generations, while reproduction biases transmission toward stronger parents, together supporting the gradual consolidation of asymmetric influence patterns in the population.

A viability constraint ensures capability never drops below a functional minimum:

$$\alpha_{t=0}^k = C^{k''} = \max(C^{k'}, 1), \quad (10)$$

where the constant 1 is a lower bound for viable capability rather than a normalization factor. We use $\alpha_{t=0}^k$ instead of C^k to highlight that the α notation represents an agent’s social or functional standing, indexed by time t . Although in the present baseline model α_t^i remains static—reflecting only inherited genetic capability—the time-dependent form anticipates future extensions where ability may evolve through experience accumulation, learning, or reputation updates.

Because initial C_i values are sampled from a normal distribution centered at 100, inherited capacities $C^{k''}$ preserve the parental scale, maintaining heterogeneity across generations and preventing collapse to unit normalization.

2.4 Speaker Selection in Cooperation

During cooperative interactions, the speaker functions as a temporary coordination focal point selected endogenously from the interacting group, rather than as a fixed or guaranteed leader. Speaker selection is therefore contingent on local interaction dynamics and reflects transient coordination demands rather than persistent authority.

To select a speaker among cooperating agents, we use a softmax-based stochastic rule weighted by ability scores α_t^i . The probability that agent i becomes the speaker at time t is

$$P_t^i(S) = \frac{\exp(\beta \alpha_t^i)}{\sum_{j=1}^n \exp(\beta \alpha_t^j)}, \quad (11)$$

where n denotes the number of cooperating agents. The parameter $\beta > 0$ controls the selection pressure: larger β increases the probability that agents with higher α_t^i are selected as speakers. This formulation favours agents with higher α_t^i while preserving stochasticity, so leadership emerges probabilistically rather than deterministically. In comparison, a linear proportional rule [20], $P_t^i(S) = \frac{\alpha_t^i}{\sum_{j=1}^n \alpha_t^j}$, underestimates small ability gaps and saturates when ability distributions are skewed. The softmax function increases sensitivity to minor disparities in α_t^i , giving slightly stronger agents a proportionally higher chance to lead. Conceptually, this corresponds to a multinomial Logit (discrete choice) model in which selection probabilities scale exponentially with perceived ability. As a result, hierarchy formation becomes smoother and continuously responsive to marginal competence differences, as reflected in the gradual rise in leadership likelihood for small $\Delta\alpha_t^i$ observed in our results.

2.5 Consensus-Based Decision-Making and Resource Allocation

During cooperative tasks, resource distribution is governed by a decentralized consensus process. Each agent i possesses an ability score α_t^i , representing task-relevant competence or influence. One agent is stochastically selected as the speaker (s), while the remaining agents ($\mathcal{L}$) act as listeners.

The group's response to the speaker is captured by a dynamic *consensus coefficient* $d_{t,q}$, which regulates alignment or resistance during the q -th cooperative event at time t :

$$d_{t,q} = 1 + \frac{1}{|\mathcal{L}|} \sum_{l \in \mathcal{L}} (\alpha_t^s - \alpha_t^l)(d_t^s - d_t^l), \quad (12)$$

where d_t^s and d_t^l denote the expected dominance of the speaker and listener, respectively, derived from each agent's normalized percentile rank $\mathcal{R}_t^i \in (0, 1]$ within the cooperating group. :

$$d_t^s = 0.5 + 0.5 \mathcal{R}_t^s, \quad d_t^l = 0.5 \mathcal{R}_t^l. \quad (13)$$

Higher ranks yield larger d_t^i , strengthening influence in the consensus process. The asymmetric ranges $[0.5, 1]$ for speakers and $[0, 0.5]$ for listeners allow both dominance and partial resistance. When the speaker is substantially stronger, $d_{t,q} > 1$ and consensus favors the speaker; when weaker, $d_{t,q} < 1$, indicating resistance. For small ability

differences, $d_{t,q}$ remains near 1, corresponding to neutral consensus. Although Eq. (12) introduces local positive feedback, hierarchical amplification is not guaranteed and remains contingent on stochastic speaker selection and broader evolutionary dynamics.

Given $d_{t,q}$, resources are allocated as:

$$r_t^i = \frac{1 + d_{t,q} \alpha_t^i}{\sum_{z \in \mathcal{T}} (1 + d_{t,q} \alpha_t^z)}. \quad (14)$$

This rule amplifies or suppresses ability differences depending on the group’s collective stance. In each cooperative cycle, $d_{t,q}$ is updated once and $\mathcal{R}_t^i$ recalculated from the current ability distribution, allowing adaptive evolution of inequality.

Table 1: Model parameters related to the consensus mechanism.

Symbol	Description	Typical Role
α_t^i	Individual ability	dynamic variable
$\mathcal{R}_t^i$	Percentile rank within team	(0,1]
d_t^s, d_t^l	Expected dominance	speaker/listener
$d_{t,q}$	Consensus coefficient	dynamic
r_t^i	Resource share	0–1

This consensus mechanism links local coordination with emergent inequality and provides a key pathway through which hierarchy stabilizes in the evolving population.

2.6 Monitoring Hierarchical Emergence via Trophic Analysis

To quantify the emergence and stability of hierarchical structure in the multi-agent system, we employ the *Trophic Incoherence* (TI) metric. Originally developed for food-web analysis, TI has been extended to characterize directed social and interaction networks [13, 14, 25, 21]. It measures how much directional links deviate from a strictly ordered hierarchy, providing a continuous index of structural coherence. Preliminary ablation checks further confirm that both mutation amplitude and initial heterogeneity are essential—disabling either prevents the system from maintaining low TI, highlighting their indispensable role in sustaining hierarchical order.

Let $W = [w_{ij}]$ denote the weighted adjacency matrix of listener-to-speaker interactions, where each element w_{ij} counts how many times agent i (as a **listener**) has chosen agent j (as a **speaker**). Each cooperative event thus represents a recommendation, so W encodes accumulated endorsements from follower to leader. Since influence flows from speaker to listener, we use the transpose W^T in all calculations.

For each agent i , we define:

$$u_i = \sum_j (w_{ij} + w_{ji}), \quad v_i = \sum_j (w_{ij} - w_{ji}), \quad (15)$$

where u_i is the total interaction strength and v_i the net bias between outgoing and incoming links. We then construct a symmetric Laplacian-like matrix $\Lambda = \text{diag}(u_i) -$

$(W + W^T)$, and solve $h = \Lambda^\dagger \cdot v$, where $h = [h_i]$ gives each agent's trophic level, $\Lambda^\dagger$ is the Moore–Penrose pseudoinverse, and $v = [v_i]$ is the interaction vector.

Trophic incoherence is then:

$$F(h) = \frac{1}{\sum_{i,j} w_{ij}} \sum_{i,j} w_{ij} (h_j - h_i - 1)^2. \quad (16)$$

Here, $F(h) \in [0, \infty)$ quantifies deviation from perfect hierarchy, where $h_j = h_i + 1$ indicates ideal downstream flow from speaker j to listener i . Lower $F(h)$ reflects more layered, directional organization; higher values imply flatter or cyclic structures. In particular, $F(h) \rightarrow 0$ denotes a perfectly stratified hierarchy, whereas $F(h) \geq 1$ marks increasing incoherence where directionality is lost. This continuous measure tracks how local interactions sustain or disrupt emergent hierarchy over time.

Self-loops are excluded; for disconnected networks, TI is computed on the largest connected component, with isolated nodes omitted from the $F(h)$ normalization. In the following, we refer to $F(h)$ as TI.

3 Results and Analysis

We first examine how the final Trophic Incoherence (TI) varies across combinations of initial heterogeneity (c) and mutation amplitude (u). Lower TI values correspond to more stratified, hierarchical states, whereas higher values indicate more disordered configurations. Figure 2 reports the mean final TI after 30 000 simulation steps, averaged over 20 independent runs, a horizon sufficient for convergence while mitigating stochastic fluctuations. As u increases, TI decreases markedly, indicating that stronger mutation facilitates the emergence of hierarchical structure. By contrast, variation in c has a weaker but non-negligible effect, most apparent at low u , suggesting that initial heterogeneity primarily influences early formation rather than long-term persistence.

3.1 Parameter Space Overview

Figure 2 summarises the mean final Trophic Incoherence (TI) across the (c, u) parameter space, where lower values indicate more stratified hierarchical organisation and higher values correspond to disordered or egalitarian states. TI is computed after 30 000 simulation steps and averaged over 20 independent runs. As mutation amplitude u increases, TI decreases sharply, indicating that stronger mutation promotes the emergence of hierarchical structure. By contrast, when mutation is minimal ($u \leq 0.1$), TI remains consistently high (> 0.9) across all c values, and no stable hierarchy forms. This regime effectively serves as an implicit ablation: in the near absence of stochastic variation, directional differentiation fails to persist, highlighting mutation as a necessary condition for hierarchy formation.

Within the explored parameter range, variation in c exerts a secondary but non-negligible influence—most visible at low u —suggesting that initial heterogeneity mainly affects the early onset rather than the long-term persistence of hierarchy. To assess robustness, we next quantify cross-run variability in the final TI. Figure 3 reports the




Fig. 2: **Mean Final Trophic Incoherence across parameter space.** Each cell shows the mean final TI across 20 replications. Darker blue tones denote lower TI, indicating stronger hierarchical order.




Fig. 3: **Variability of TI across replications (IQR).** The interquartile range captures cross-run variability for each (c, u) pair. Smaller values indicate greater stability and convergence of hierarchical order.

interquartile range (IQR) for each (c, u) pair. Smaller IQR values indicate higher stability and convergence of hierarchy formation.

Large mutation amplitudes ($u = 1$ or $u = 2$) produce both low mean TI (Fig. 2) and small IQR, implying that hierarchies formed under these conditions are not only stronger but also highly reproducible. In contrast, intermediate mutation levels ($u \approx 0.5$) show the highest variability, revealing a transition zone where outcomes fluctuate between ordered and disordered regimes. Here, reproducibility refers to *cross-run consistency*

rather than temporal invariance, with $\text{IQR}(TI) < 0.05$ indicating consistent outcomes across independent runs.




Fig. 4: **Phase map of hierarchical regimes.** Regions where $\text{median}(TI) < 0.45$ and $\text{IQR} < 0.05$ are classified as Consistent decrease (dark blue), indicating strong cross-run reproducibility across independent runs; transitional regions as Rebound (orange), and disordered regions as No-change (gray). The map highlights a clear boundary separating ordered and disordered regimes.

Combining these two indicators, we identify distinct dynamical regimes in the (c, u) space. Figure 4 shows how these regimes vary across 20 independent runs. As mutation amplitude increases, cross-run consistent trajectories become dominant, while rebound and no-change outcomes diminish. At intermediate u values (≈ 0.5), all three regimes coexist, marking a probabilistic transition between ordered and disordered behavior. Together, these three maps (Figs. 2–4) reveal a phase-like transition across the parameter space: low mutation maintains disordered, high-TI states; intermediate u leads to mixed regimes; and high u consistently produces low-TI, stable hierarchical order.

3.2 Temporal Dynamics of Hierarchy

To examine not only *whether* hierarchy emerges but also *how* it develops, we trace the time series of Trophic Incoherence (TI) over the whole simulation.

Figure 5 reports a representative case ($c = 0.05$, $u = 1.0$). At the beginning, TI rises quickly to a high level as interaction links are still noisy and weakly ordered. After this transient, TI starts to decrease monotonically and eventually fluctuates within a narrow band. The blue curve shows that, on average, the system converges toward a low-TI configuration, while the gray trajectories reveal that individual runs may reach this regime at slightly different times. The green shaded segment marks the period in which both criteria $\text{median}(TI) < 0.45$ and $\text{IQR} < 0.05$ are satisfied, i.e. when a stable hierarchical regime has effectively formed.




Fig. 5: **Representative trajectories at ($c = 0.05$, $u = 1.0$)**. Each gray line represents one of 20 runs; the blue curve is the median TI, and the shaded area the interquartile range (25–75%). The green region highlights periods satisfying the stability criterion ($\text{median}(TI) < 0.45$ and $\text{IQR} < 0.05$), indicating the onset of a stable hierarchical regime.

To see how robust this pattern is across parameter settings, we further compare TI trajectories under different mutation amplitudes u and initial heterogeneity c (Fig. 6). A common structure appears in all panels: an early exploration phase with high TI, followed by an ordering phase in which TI declines, and finally a late stage in which TI either stabilizes or rebounds slightly. However, the *speed* and *depth* of this decline depend on u . When mutation is weak ($u = 0.1$), TI remains high for the whole run, showing that almost no durable hierarchy can be sustained. For intermediate values




Fig. 6: **Temporal dynamics of TI under varying mutation strengths u** . Each curve shows the median TI across 20 replicates (shaded area = IQR 25–75%); rows correspond to different initial heterogeneity c . Higher u values lead to earlier and deeper ordering, whereas low u maintains high-TI, weakly ordered states.

($u = 0.2$ or $u = 0.5$), TI decreases but the spread between runs widens, producing the mixed and “transitional” regimes already identified in Fig. 4. When mutation is strong ($u = 1.0$), TI drops earliest and reaches the lowest plateau, indicating that continuous trait variation actually helps the system select and reinforce a directional interaction structure. Across different c , this temporal profile is largely preserved, confirming that ongoing mutation is a stronger driver of hierarchical ordering than the initial spread of abilities.

In combination with the parameter-space results, these temporal curves show that hierarchy in this model is not created instantaneously but is gradually consolidated as successful speaker–listener relations accumulate in the recommendation network. Mutation that is too weak does not provide enough variation for such relations to be filtered; mutation at moderate to high levels accelerates this consolidation and keeps the system in the low-TI band.

3.3 Structural Interpretation through Network Comparison

To illustrate the structural meaning of Trophic Incoherence (TI), we visualize two representative communication networks obtained at the end of the simulation. Each network is arranged using a *trophic-level layout* that (i) places lower-level agents near the bottom and (ii) minimizes edge crossings. When hierarchy is well defined, these criteria yield clear stratified layers; when it is weak, crossings cannot be minimized without breaking vertical order, producing a tangled and symmetric pattern. Thus, the visual contrast between the two panels directly reflects the presence or absence of hierarchical organization rather than mere node rearrangement.

Each node in Fig. 7 represents a surviving agent positioned vertically by its trophic level h_i . Directed edges indicate accumulated listener-to-speaker recommendations,




Fig. 7: Network Topologies at High and Low TI. (a) Communication network at step 30 000 with $F(h) \approx 0.99$; (b) at step 30 000 with $F(h) \approx 0.31$. Both are visualized using a trophic-level layout positioning lower-level nodes at the bottom while minimizing edge crossings. Strong hierarchical ordering, as in (b), forms distinct tiers, whereas incoherent networks like (a) appear tangled and symmetric.

while node size and color encode speaking frequency. Larger, darker nodes thus denote agents with higher social endorsement, integrating influence and hierarchical position within the network.

The right-hand network ($F(h) \approx 0.31$) shows clear trophic layering: agents cluster into distinct tiers, and a few dominant nodes receive disproportionate endorsement, reflecting persistent leadership and asymmetric information flow. By contrast, the left-hand network ($F(h) \approx 0.99$) appears diffuse and symmetric, lacking both tiers and concentration of influence—an egalitarian configuration without sustained dominance. These contrasts illustrate that low TI coincides with visually stratified networks, whereas high TI is associated with diffuse, weakly ordered structures.

4 Discussion and Future Work

This study investigated how hierarchical structure can arise in a decentralised multi-agent system without predefined leaders. In our model, agents coordinate through local foraging and speaker–listener interactions, and these interactions gradually induce asymmetric influence patterns captured by the Trophic Incoherence (TI) index.

Across the parameter space, mutation amplitude u plays the dominant role in determining whether hierarchy forms. When mutation is weak, inherited traits remain nearly static and the population stays in a disordered, high-TI regime. As u increases, heritable variation is continuously reintroduced, enabling directional relationships to be filtered and reinforced across generations. This produces a clear transition around $u \approx 0.5$ – 1.0 , where TI drops and stable hierarchy becomes the most likely outcome. Initial heterogeneity c mainly affects the early phase of convergence: higher c delays ordering but does not prevent low-TI states once reinforcement accumulates.

The temporal profiles of TI clarify this mechanism. Under moderate u , TI declines steadily as cooperative interactions repeatedly endorse particular speaker–listener pairs, while network visualisations show that low-TI runs develop clear trophic layering and concentrated influence, whereas high-TI runs remain diffuse and symmetric. TI therefore provides a practical link between local coordination dynamics and population-level organisation in the model.

Although each agent’s capability α^i is fixed in the current implementation, its time-indexed notation α_t^i makes the framework compatible with extensions in which ability evolves through learning, reinforcement, or social experience. The present model intentionally prioritises structural hierarchy formation over task-level performance or demographic richness; extensions such as adaptive foraging efficiency, resilience to agent loss, and more diverse gender representations are left for future work.

In summary, hierarchy in this system arises from the interaction of variation, competition, and reinforcement rather than explicit design. Future work will extend the model by introducing experience-based updates to α_t^i , allowing environmental conditions to fluctuate, and examining resilience under deception or rapid population turnover. The source code for the simulation model and analysis scripts is publicly available at: <https://github.com/ShanshanMao999/cooperation-to-hierarchy-abm>.

References

1. Boehm, C.: *Hierarchy in the Forest: The Evolution of Egalitarian Behavior*. Harvard University Press (2009)
2. Couzin, I.D.: Collective cognition in animal groups. *Trends in Cognitive Sciences* **13**(1), 36–43 (2009)
3. Cristelli, M., Batty, M., Pietronero, L.: There is more than a power law in zipf. *Scientific reports* **2**(1), 812 (2012)
4. Daniel, P.A., Daniel, E.: Multi-level project organizing: A complex adaptive systems perspective. In: *Research handbook on complex project organizing*, pp. 138–147. Edward Elgar Publishing (2023)
5. De Marzo, G., Gabrielli, A., Zaccaria, A., Pietronero, L.: Dynamical approach to zipf’s law. *Physical Review Research* **3**(1), 013084 (2021)
6. Detrain, C., Deneubourg, J.L.: Collective decision-making and foraging patterns in ants and honeybees. *Advances in Insect Physiology* **35**, 123–173 (2008)
7. Dridi, S., Akçay, E.: Learning to cooperate: The evolution of social rewards in repeated interactions. *The American Naturalist* **191**(1), 58–73 (2018)
8. Foss, N.J., Schmidt, J., Teece, D.J.: Ecosystem leadership as a dynamic capability. *Long range planning* **56**(1), 102270 (2023)
9. Gagnon, S.A., Brunyé, T.T., Robin, C., Mahoney, C.R., Taylor, H.A.: High and mighty: Implicit associations between space and social status. *Frontiers in Psychology* **2**, 259 (2011)
10. Gotelli, N.J.: *A Primer of Ecology*. Sinauer Associates (2008)
11. Grimaldo, F., Lozano, M., Barber, F.: Madem: a multi-modal decision making for social mas. In: *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems (AAMAS)*. vol. 1, pp. 183–190. IFAAMAS (2008)
12. Haertel, T.M., Leckelt, M., Grosz, M.P., Kuefner, A.C., Geukes, K., Back, M.D.: Pathways from narcissism to leadership emergence in social groups. *European Journal of Personality* **37**(1), 72–94 (2023)
13. Johnson, S., Domínguez-García, V., Donetti, L., Muñoz, M.A.: Trophic coherence determines food-web stability. *Proceedings of the National Academy of Sciences* **111**(50), 17923–17928 (2014)
14. Klaise, J., Johnson, S.: From neurons to epidemics: How trophic coherence affects spreading processes. *Chaos: An Interdisciplinary Journal of Nonlinear Science* **26**(6), 065310 (2016)
15. Lane, D.: Hierarchy, complexity, society. In: *Hierarchy in natural and social sciences*, pp. 81–119. Springer, Dordrecht (2006)
16. Macal, C.M., North, M.J.: Tutorial on agent-based modeling and simulation. In: *Proceedings of the Winter Simulation Conference, 2005*. pp. 14–pp. IEEE (2005)
17. Migliano, A.B., Vinicius, L.: The origins of human cumulative culture: from the foraging niche to collective intelligence. *Philosophical Transactions of the Royal Society B* **377**(1843), 20200317 (2022)
18. Papasimeon, M.: *Modelling agent-environment interaction in multi-agent simulations with affordances*. Citeseer (2009)
19. Perret, C., Hart, E., Powers, S.T.: From disorganized equality to efficient hierarchy: how group size drives the evolution of hierarchy in human societies. *Proceedings of the Royal Society B* **287**(1928), 20200693 (2020)
20. Perret, C., Powers, S.T., Hart, E.: Emergence of hierarchy from the evolution of individual influence in an agent-based model. In: *Artificial Life Conference Proceedings*. pp. 348–355. MIT Press (2017)
21. Pilgrim, C., Guo, W., Johnson, S.: Organisational social influence on directed hierarchical graphs, from tyranny to anarchy. *Scientific Reports* **10**(1), 4388 (2020)

22. Redhead, D., Cheng, J.T., Driver, C., Foulsham, T., O’Gorman, R.: On the dynamics of social hierarchy: A longitudinal investigation of the rise and fall of prestige, dominance, and social rank in naturalistic task groups. *Evolution and Human Behavior* **40**(2), 222–234 (2019)
23. Redhead, D., Power, E.A.: Social hierarchies and social networks in humans. *Philosophical Transactions of the Royal Society B* **377**(1845), 20200440 (2022)
24. Ren, T., Zeng, X.J.: Enhancing cooperation through selective interaction and long-term experiences in multi-agent reinforcement learning. *arXiv preprint arXiv:2405.02654* (2024)
25. Rodgers, N., Tiño, P., Johnson, S.: Network hierarchy and pattern recovery in directed sparse hopfield networks. *Physical Review E* **105**(6), 064304 (2022)
26. Sayama, H.: Introduction to the modeling and analysis of complex systems. Open SUNY Textbooks (2015)
27. Tibbetts, E.A., Pardo-Sanchez, J., Weise, C.: The establishment and maintenance of dominance hierarchies. *Philosophical Transactions of the Royal Society B* **377**(1845), 20200450 (2022)
28. de Weerd, M.M., Zhang, Y., Klos, T.: Multiagent task allocation in social networks. *Autonomous Agents and Multi-Agent Systems* **25**(1), 46–86 (2012)
29. Zefferman, M.R.: Constraints on cooperation shape hierarchical versus distributed structure in human groups. *Scientific Reports* **13**(1), 1160 (2023)
30. Łomnicki, A.: Individual-based models and the individual-based approach to population ecology. *Ecological Modelling* **115**(2-3), 191–198 (1999)