Collective Behavior Emerging from Social Learning Strategies and Network Structures

Jingyu G. Xi (jingyu.xi@uni-konstanz.de)
The Centre for the Advanced Study of Collective Behaviour, University of Konstanz
Universitätsstraße 10, Konstanz, 78464 Germany

Wataru Toyokawa (wataru.toyokawa@riken.jp)
Computational Group Dynamics Collaboration Unit, BTCC, Riken CBS
2-1 Hirosawa Wako City, Saitama 351-0198 Japan

Abstract

Humans make decisions collectively by combining individual and social learning. Individuals benefit from groups when individual exploration fails to accurately assess the environment, a phenomenon known as the "wisdom of crowds." Previous studies indicate that self-organizing group dynamics can reduce suboptimal biases in noisy environments, particularly in fully connected groups. However, agents often have only partial information due to cognitive and physical constraints. To explore how diverse social network structures influence the collective dynamics of social learners, we integrate a decentralized network with a social reinforcement learning model in repeated two-armed bandit tasks. Our results suggest that: 1. Social learning in a sparse network outperforms asocial solo learning in highly uncertain tasks. 2. The phenomenon "less is more, and more is different" holds true only when agents strategically balance individual and social learning. 3. The group size effect on collective performance is significantly influenced by network structures.

Keywords: Collective Behavior; Social Reinforcement Learning; Network Dynamics; Risk Aversion

Introduction

Understanding the principles of collective behavior has been an enduring conundrum in behavioral sciences. Previous work suggests that conformist social learning, whereby the most common behavior in a group is disproportionately more likely to be copied, can regulate both the wisdom and "madness" of interactive human groups (Toyokawa & Gaissmaier, 2022; Toyokawa et al., 2019). By combining reinforcement learning with a decision-biasing imitation process (Najar et al., 2019), agent-based model simulations have identified a critical level of conformity bias, above which collective behavior becomes dominated by herding, resulting in inflexible collective decision-making. On the other hand, collective learners can remain flexible as long as reliance on the social information remains sufficiently low (Lorenz et al., 2011). These studies suggest that the conflict between collective intelligence and maladaptive herding could potentially be predicted with quantitative knowledge of the social learning strategies employed.

However, past studies on human collective learning have been limited to circumstances in which individuals could observe all other individuals in the same group (Collet et al., 2023; Deffner et al., 2024; Shirado, 2022). In other words, our understanding thus far has been based on fully connected social networks. Given that the source of social information is often restricted to socially connected individuals, and that non-random social network structures have been shown to be important in understanding human collective behavior in real societies (Watts & Strogatz, 1998), it is necessary to synthesize our understanding of how social learning strategies deployed by individuals interact with population-level social network architecture to shape collective dynamics.

Previous studies on cooperation (Rand et al., 2014; Shirado & Christakis, 2020), social foraging (Cantor & Farine, 2018), coordination (Angus & Newton, 2020; Shirado & Christakis, 2017), collective intelligence and network structures in NK problem-solving tasks (Lazer & Friedman, 2007), and cultural evolution (Cantor et al., 2021; Smolla & Akçay, 2019) have shown that information transfer across structured social connections can qualitatively alter the dynamics of collective behavior and its evolution, compared to collective behavior observed in random or fully connected social networks. By exploring a range of network connectivities, previous work has generally concluded that the denser the network, the higher the risk of herding due to increased correlations of opinions between neighboring agents (Dasaratha & He, 2020; Reina et al., 2024). Similarly to observations in conformist social learning studies, the conflict between collective intelligence and maladaptive herding appears to be predictable with quantitative knowledge of the realized social network structure.

However, the decision-making processes assumed in these social network studies were often overly simplified, where individuals merely made a one-shot decision based on provided social cues and, therefore, did not actively engage in self-evaluated trial-and-error individual learning during social interactions (Brackbill & Centola, 2020; Das et al., 2014; Lazer & Friedman, 2007). For example, in collective behavior studies with the NK model, there is usually no learning per se, but instead more emphasis on search and selection (Lazer & Friedman, 2007). As a result, it remains unclear how diverse social structures interact with the learning strategies deployed. Since a lack of individual exploration often causes social learners to become stuck in suboptimal decision-making within fully connected networks (Toyokawa & Gaissmaier, 2022), it is possible that certain types of less-connected networks could "rescue" clusters of conformist learners by facilitating the flow of new information from other parts of the network. On the other hand, the risk of herding in dense networks might be mitigated by the presence of more explorative individuals.

The key questions we aim to address here are threefold:

In D. Barner, N.R. Bramley, A. Ruggeri and C.M. Walker (Eds.), *Proceedings of the 47th Annual Conference of the Cognitive Science Society*

- How does collective decision-making performance, compared to both solo learning and fully connected social learning, benefit from less dense social connections?
- What kinds of nonlinear collective decision phenomena emerge in both risky and non-risky decision environments?
- How does individual-level decision performance influence emergent collective behavior?

Computational Models

The Social Two-Armed Bandit Task

To address these questions, we employed a social "two-armed bandit" task (Figure 1A), in which individuals could observe others' decision-making behavior (social frequency information) but not their payoff information. Unlike the previous "fully-connected" studies (Toyokawa & Gaissmaier, 2022; Toyokawa et al., 2019), social information was available only from predetermined social connections (Figure 1A; Toyokawa et al., 2019). To manipulate individual exploration and exploitation patterns, and to potentially impact the resulting collective behavior (Toyokawa & Gaissmaier, 2022), we tested two different reward distributions (Figure 1B). In the risk-homogeneous setting, the variance of payoffs from both behavioral options was equal, whereas in the risk-heterogeneous setting, the variance of payoffs from one option (the riskier option) was higher than that of the other option (the safer option). We set the expected payoff of the riskier option to always be higher than that of the safer option. It is known that individual trial-and-error learning tends to result in avoidance of such a riskier option, even if it provides higher expected payoffs in the long run, due to biased exploration (a phenomenon known as the hot stove effect; Denrell, 2007, 2024; Denrell and March, 2001). Therefore, the risk-heterogeneous setting allowed us to examine how such biased decision-making operating at the individual level might interact with both social learning and network structure.

Network Manipulations

We manipulated network density using the "2m-ring" network algorithm to explore the importance of a more distributed system without central individuals (or "nodes" of the social network). The 2m-ring network algorithm generates networks in which each node is connected to its m nearest neighbors on both the "left-hand side" and the "right-hand side" (Figure 1A). This approach ensures that all nodes have the same number of connections. Denoting the number of connections each node has as $\langle k \rangle$, which is called "degree", then each node has degree $\langle k \rangle = 2m$. We assumed that the network structure does not change over the course of the bandit task, and that social information transmits in both directions between connected nodes. Denoting the number of connections ("edges") in the network as $|E|$ and the number of nodes ("vertices") as $|V|$, the density of a given social network is defined as the ratio of the number of existing edges to the maximum possible number of edges, given by D_G:

$$D_G = \frac{2|E|}{|V|(|V|-1)}. \tag{1}$$

Past results from studies on network dynamics and the wisdom of crowds (Becker et al., 2017) demonstrate that decentralized systems, where individuals have equal social influences, result in higher accuracy in a group task compared to centralized networks. Popular network models such as the scale-free model (Li et al., 2007) and the random network model (Barabâsi et al., 2002) that have been widely tested in network science and information diffusion studies, are great abstractions representing real-world complex networks. However, due to the randomness and noise those models naturally possess, the results might be difficult to disentangle and interpret if the networks generated by these algorithms are centralized. Thus our work uses the more decentralized "2m-ring" network to explore the mechanism between collective learning and network features.

We also manipulated group size defined as the number of nodes $|V|$ to examine its effect on collective performance in social learning. Some studies have found that large groups outperform smaller ones in complex tasks (Mao et al., 2016) or spatial discrimination tasks (Langley et al., 2018), while other studies have shown that smaller groups perform better (Kao & Couzin, 2014; Liden et al., 2004). The reasons behind these inconsistent group size effects remain controversial. In our study, we included group size as a network parameter and investigated its interaction with other network and learning parameters.

Baseline Reinforcement Learning

For the value learning process at the individual level, we assumed that agents update $Q_{i,t}(a)$, a belief of option a's value at time t for agent i, according to the following forms:

$$Q_{i,t+1}(a) = \begin{cases} Q_{i,t}(a) + \alpha\left(R_{i,t}(a) - Q_{i,t}(a)\right), & \text{if } a_{i,t} = a \text{ or } t = 1 \\ Q_{i,t}(a), & \text{otherwise} \end{cases} \tag{2}$$

where $\alpha \in [0,1]$ is a constant learning rate parameter that controls the step size of value updating. $Q_{i,t}(a)$ is updated only when option a is selected at time t.

Then, agents choose an option based on the Q-values mapped onto the softmax function as follows:

$$P_{i,t+1}(a)^{asocial} = \frac{exp(\beta Q_{i,t}(a))}{\sum_m exp(\beta Q_{i,t}(a))} \tag{3}$$

where m is the total number of options and $\beta \in [0,\infty)$ is the inverse temperature parameter that controls the stochasticity during the action selection process, which impacts the exploration-exploitation trade-off. Agents with low β values tend to make more exploratory decisions, with actions being uniformly at random when $\beta = 0$, while agents with high β tend to select actions associated with higher Q-values, with actions becoming deterministic as $\beta \to \infty$.

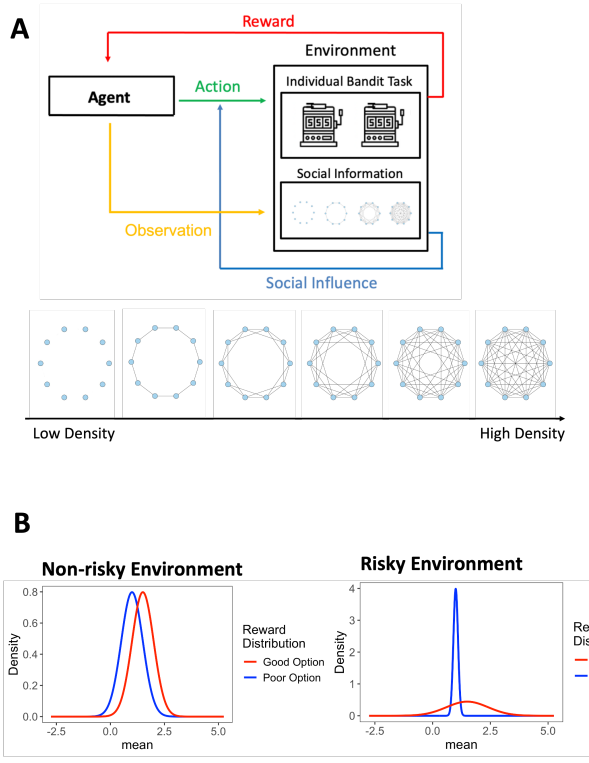

Figure 1: Modal overview. **A**, Social MAB with static 2m-ring network structures. **B**, Reward distributions in two types of environments. Non-risky: $A_{poor} \sim N(1, 0.25)$, $A_{good} \sim N(1.5, 0.25)$. Risky: $A_{safe} \sim N(1, 0.01)$, $A_{risky} \sim N(1.5, 0.81)$.

Conformist Social Learning

For the social learning process, we assumed social frequency-dependent decision-making (that is, a decision-biasing imitation process; Najar et al., 2019) in which the social information influences the action probability $P_{i,t+1}(a)$ as a linear combination of both the Q-value-based softmax process $P_{i,t+1}(a)^{\text{asocial}}$ and the conformist social influence $P_{i,t+1}(a)^{\text{social}}$, given by:

$$P_{i,t+1}(a) = (1-\sigma)P_{i,t+1}(a)^{asocial} + \sigma P_{i,t+1}(a)^{social}, \quad (4)$$

where the combination is weighted by σ ($0 < \sigma < 1$), the social learning weight. For the social influence term $P_{i,t+1}(a)^{\text{social}}$, we assumed a conformity function:

$$P_{i,t+1}(a)^{social} = \frac{N_{i,t}(a)^{\theta}}{\sum_k N_{i,t}(m)^{\theta}} \quad (5)$$

where $N_{i,t}(a)$ is the number of connected agents (from the perspective of individual i) who chose option a at time t, and $\theta \in (-\infty, \infty)$, called the conformity exponent, controls the strength of conformity bias—that is, how strongly agents conform to the majority behavior out of m options. When $\theta > 1$, agents are disproportionately more likely to follow the majority's decision-making, while $\theta = 1$ yields a proportional response to the social frequency. When $0 < \theta < 1$, agents develop anti-conformity, favoring the minority option, while $\theta = 0$ implies no social influence and agents act uniformly at random.

Agent-based Model Simulations

We ran agent-based model simulations by systematically varying the learning parameters as well as the social network structure. Unless otherwise stated, the bandit task setup was the same as that shown in Figure 4.

Results

The simulations identified three distinct patterns of emerging collective behavioral dynamics, irrespective of the network structures employed (Figure 2B). First, there was a parameter combination in which groups of agents exhibited a single collective attractor, resulting in a unimodal distribution of decision-making performance whose mode was biased toward the high-payoff option (henceforth referred to as "collective intelligence"). Second, when both the conformity bias (θ) and the social learning weight (σ) were large, groups tended to show bimodal distributions, whereby some groups exploited the higher-payoff option while others repeatedly chose the lower-payoff option due to strong positive feedback (referred to as "herding"). Third, it was also possible for groups to exhibit a unimodal distribution centered on the low-payoff option.

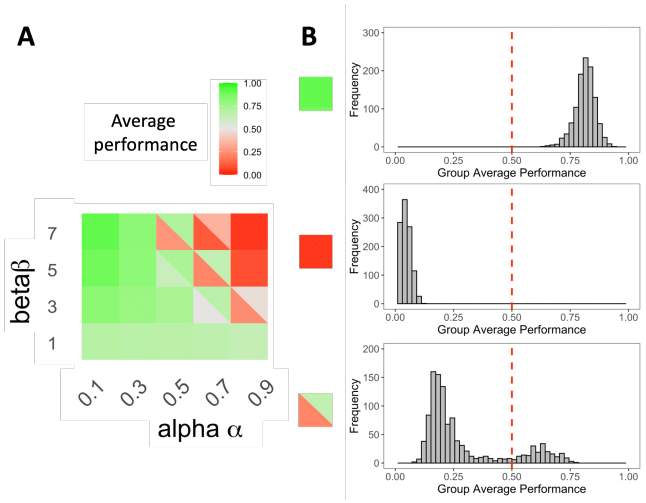

Figure 2: **A**, Demonstration of the mean choice probability over the last 10 time steps with fixed σ and θ, after fitting the original distributions from 100 repetitions (shown below) using a Gaussian Mixture Model. $p = 0.5$ represents chance level in our setup. **B**, Original choice probability distributions from the last 10 time steps, demonstrating three types of collective dynamics.

Collective performance under sparse connections

We first examined the relationship between network density and collective decision performance. Figure 3A shows the average performance for an empty graph of 50 nodes (i.e., 50 solo learners), $G_{50,\langle k\rangle=0}$. In the risk-homogeneous environmental setting (Figure 3A, left), all agents perform rela-

tively well regardless of the learning parameters. In contrast, in the risk-heterogeneous environment, agents may be biased toward the safer but payoff-minimizing option when α and/or β is large—a phenomenon known as the "hot stove effect" (Figure 3A, right; Toyokawa and Gaissmaier, 2022).

Then we explored the effect of social learning by adding connections, starting from a minimum ring topology, $G_{50,\langle k\rangle=2}$. In this minimum ring network, we found that social influence had opposite effects in the two risk settings. In the risk-homogeneous task, social learners (i.e., $\sigma > 0$) performed worse than solo asocial learners (Figure 3B, top left). In contrast, in the risk-heterogeneous task, this minimal level of social connection helped mitigate suboptimal risk aversion (Figure 3B, top right).

A sparse network with minimal social information can be either beneficial or detrimental to collective performance, depending on whether promoting exploration is advantageous. The addition of the social term $P_{i,t+1}(a)^{\text{social}}$ typically increases exploratory actions due to sparse social sampling. Since the balance between exploration and exploitation is key in such a bandit task, consistently increasing exploration can be detrimental—especially in the risk-homogeneous case. However, in the risk-heterogeneous task, where solo agents tend to get stuck in the suboptimal safer option, promoting exploration can help them escape the "hot stove" effect (Toyokawa & Gaissmaier, 2022).

Increasing density: sparse information facilitates the emergence of wisdom of crowds conditionally

We then examined whether increasing network density changes collective dynamics by increasing the average degree from $G_{50,\langle k\rangle=0}$ to $G_{50,\langle k\rangle=49}$. Figure 4B shows the resulting dynamics across varying densities, given $G_{50,\langle k\rangle} \in 0, 2, 12, 22, 49$. Figure 4A provides a close-up of the case $G_{50,\langle k\rangle=49}$ to display the inner panels more clearly.

We first controlled the individual learning parameters and varied the density parameter at the system level. In both risk-homogeneous and risk-heterogeneous environments, increasing network density from minimum sparsity ($G_{50,\langle k\rangle=2}$) to higher sparsity ($G_{50,\langle k\rangle=12}$) resulted in the gradual emergence of bifurcation, especially when agents had high learning parameter values (α, β, σ, θ). A sparse network mitigates suboptimal herding behavior only under certain conditions—specifically, when agents are less susceptible to estimation error, more explorative, and rely less on social information in the risky environment.

We then controlled the density parameter at the system level and varied the individual learning parameters. Increasing the learning rate (α), inverse temperature (β), and conformity exponent (θ) had effects similar to increasing network density: as shown in Figure 4B, bifurcation gradually dominates the system as α, β, and θ increase. However, the role of the social learning weight (σ) is more complex in the risky environment (Figure 4B, bottom row). Increasing σ initially mitigates risk-averse behavior but eventually contributes to suboptimal herding as σ continues to increase—see $\theta = 4$ in

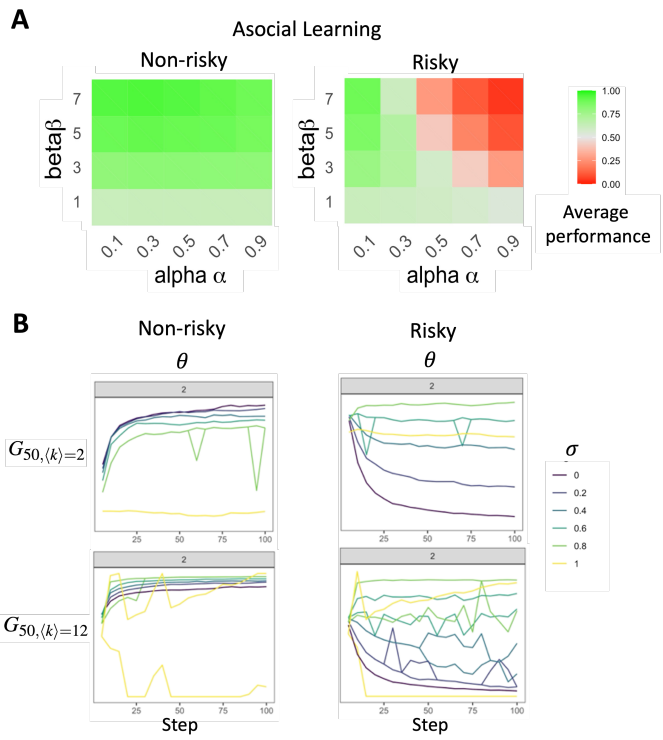

Figure 3: Simulated model performance over 100 replications.**A**, Asocial learning results: the average choice probability of the option with the highest accumulated expected value, based on fitting the distribution of the last 10 steps to a Gaussian Mixture Model (GMM) with components $K \in 1, 2$. **B**, The average choice probability plotted over time steps, given $G_{50,\langle k\rangle=2}$ and $G_{50,\langle k\rangle=12}$. When $\sigma = 0$, agents engage in asocial solo learning.

$G_{50,\langle k\rangle=12}$. The red region first decreases and then gradually transitions into a dual-colored mixture.

There is a similar functional relationship between individual-level learning parameters and system-level network parameters in both non-risky and risky environments. Increasing social conformity bias (θ) at the individual level has a similar effect to increasing network density at the collective level. However, increasing the social learning weight (σ) at the individual level functions somewhat differently, due to risk-averse behavior elicited by the hot stove effect.

Sparse networks, where agents have less social information, can often be more efficient and resilient than dense networks. However, if individuals overreact to estimation error, overexploit, and rely excessively on social information, the benefits of a sparse network diminish, and the system exhibits high volatility. "Less is more" applies only when agents do not rely too heavily on social information while retaining a certain level of individual learning.

Contradicting group size effect on the collective level and the individual level

To demonstrate how group size influences the trend of bifurcation, we varied group sizes in a risky environment while controlling and manipulating graph density at the collective

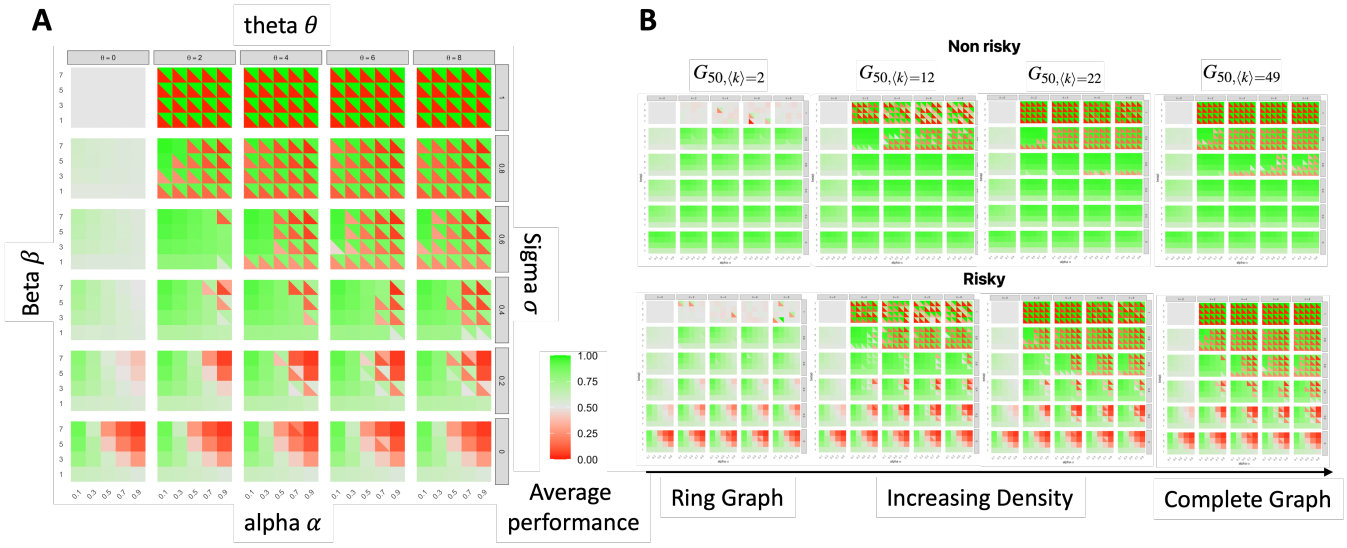

Figure 4: **A**, Demonstration of Figure 4B. Heat map of the average choice probability in the risky setup, given a complete graph $G_{50,\langle k\rangle=49}$. In each inner grid of the outer 6×5 panel, the x-axis represents the learning rate α, and the y-axis represents the inverse temperature β. Moving left to right across the outer 6×5 grid, each facet shows increasing conformity bias $\theta \in 0, 2, 4, 6, 8$; from bottom to top, each facet corresponds to increasing social learning weight $\sigma \in 0, 0.2, 0.4, 0.6, 0.8, 1.0$. **B**, From left to right: average choice probability with increasing network density, given $G_{50,\langle k\rangle} \in 0, 2, 12, 22, 49$. $G_{50,\langle k\rangle=0}$ corresponds to the case where $\sigma = 0$.

level and average degree at the individual level. In our study, group size refers to the number of individuals in the population, and average degree refers to the average number of connections each individual has. Figure 5 shows the proportion of parameter combinations that resulted in bimodal distributions of decision performance (i.e., herding) out of a total of 20 grids in the reinforcement learning parameter space, plotted as a function of group size. The results suggest that the effects of group size differ between the collective and individual levels. In Figure 5A, when group size increases while keeping network density constant, the system becomes more likely to exhibit bimodal distributions in decision performance. In contrast, in Figure 5B, when group size increases while maintaining a fixed average degree, bifurcation diminishes, and unimodal collective intelligence gradually emerges. The emergence of collective intelligence in this case can be explained by the decreasing network density that results when group size increases under a fixed degree (i.e., the network becomes sparser). Notably, the system with a lower average degree ($\langle k\rangle = 10$) outperforms the system with a higher average degree ($\langle k\rangle = 20$). Collective performance is significantly shaped by both the absolute quantity of information in the system and the social information constraints at the individual level:

- With the same relative quantity of information (i.e., fixed density), larger group sizes amplify bifurcation by increasing individual sampling size (Figure 5A).

- Large group sizes enhance the collective performance of individuals with limited information retrieval and processing capacity, but do not necessarily benefit agents with greater information capacity (Figures 5A, 5B).

Large group sizes only facilitate performance when the network is relatively sparse. A large system with overwhelming social information is more susceptible to bifurcated fluctuations, whereas smaller systems are more resilient to suboptimal dynamics, making bimodal distributions less likely to persist. In game theory and social dynamics, small groups tend to retain diverse behaviors or strategies, while larger groups are more prone to converging on a single equilibrium—a phenomenon described as "more is different" (Ghavasieh & De Domenico, 2024). However, our results suggest that this outcome can be reversed by modifying network structures (Figure 5A).

More than the sum of their parts

The relationship between group size and the likelihood of collective intelligence (i.e., increasing the chance of choosing the higher-payoff option while decreasing the risk of herding) appears to follow a nonlinear interaction. In Figure 5B, the yield of collective intelligence at the system level is not proportional to the number of system components—namely, the group size. Given $\langle k\rangle = 20$, the bifurcated distribution remains relatively robust at smaller group sizes, where density is higher, but it rapidly declines as density decreases, until it is fully mitigated in larger group sizes. The lower the average degree, the more rapidly collective intelligence emerges (i.e., the bimodal distribution diminishes) as group size increases. This trend is evidenced by the convex shape of the $\langle k\rangle = 10$ curve and the concave shape of the $\langle k\rangle = 20$ curve.

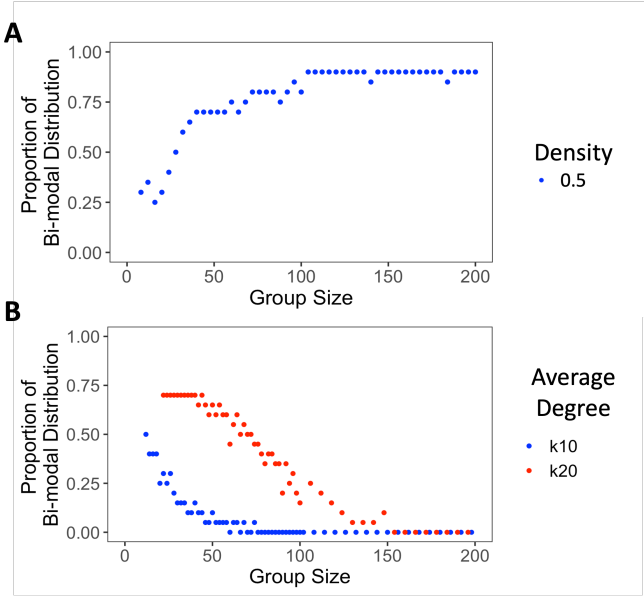

Figure 5: Proportion of occurrences of bimodal distributions as a function of group size in a risky environment. Reinforcement learning parameter space: $\alpha \in 0.1, 0.3, 0.5, 0.7, 0.9$, $\beta \in 1, 3, 5, 7$, $\sigma = 0.6$, $\theta = 6$. **A**, Fixing the density at $D = 0.5$, the average degree increases with group size, meaning that each agent has an increasing number of edges across groups. **B**, Fixing the average degree at $\langle k \rangle = 10$ and $\langle k \rangle = 20$, the network density decreases as group size increases. In this case, each agent has the same number of edges across group sizes.

Discussion

In this work, we integrated social reinforcement learning with network features to explore how the exploration–exploitation trade-off in individual learning, combined with frequency-dependent social learning strategies at the individual level and information structure at the collective level, shapes collective learning dynamics. Our simulations demonstrate that groups generally benefit from the "wisdom of crowds" when the environment is highly uncertain and risky. We also showed that collective dynamics can benefit from reduced social information in both non-risky and risky environments, even though animals—including humans—tend to actively seek social information when uncertain.

Furthermore, we highlight the similar functional roles of learning parameters at the individual level and network features at the system level, with potential implications for decision-making interventions that operate across scales. This similarity can be explained by a form of positive feedback that emerges at both levels. At the individual level, increasing the values of social learning parameters reduces the benefits of stochastic exploration and amplifies the influence of social information. At the system level, increasing network density gives each agent access to more social informants, leading to more accurate sampling of social information across the population. If agents happen to discover the optimal solution early, the system can rapidly converge on it. However, if agents become stuck in suboptimal options, high social learning values and dense connectivity hinder the stochastic variability that might otherwise help the group escape from this pitfall.

In addition to the similar cross-scale functions, we also found a nonlinear relationship between collective performance and group size. This type of nonlinear relationship between system-level outcomes and individual components has also been observed in biological systems, such as in ants. In ant collective foraging behavior, the increase in colony size is not proportional to the number of ants at the feeder, due to a positive reinforcement effect that emerges in larger colonies but not in smaller ones (Beekman et al., 2001; Sumpter, 2006). Our results suggest that group structure, in addition to group size, contributes to this nonlinearity in partially connected networks.

This study marks the beginning of an exploration into how information structures influence the emergence of collective intelligence from self-organizing individual behavior, using a reinforcement learning framework. Many network features that characterize real-world social interactions remain to be examined. Nevertheless, this study represents one of the first attempts to integrate cross-scale parameters into learning models and reveals novel patterns that were not apparent in existing theories.

Limitations and Future Directions

One limitation that warrants attention is that all network structures tested in our study are distributed, static topologies with fixed edges over time. The equally distributed system—where each agent has exactly the same amount of social information—does not account for the heterogeneous cognitive or physical constraints of animals. Additionally, our model simplifies social interactions by enforcing an undirected graph in the multi-armed bandit problem, thereby missing the opportunity to explore unidirectional social learning (i.e., from observers to demonstrators).

Beyond network structure limitations, our treatment of risk-averse behavior assumes it is detrimental to the goal of maximizing accumulated rewards. However, risk aversion does not always lead to negative outcomes and may offer adaptive benefits in specific contexts.

In future work, we plan to explore how evolving, dynamic networks shape collective behavior, particularly in settings where agents exhibit agency by actively selecting and rewiring their social connections. In summary, our integration of social reinforcement learning and network features deepens our understanding of the relationship between individual and collective intelligence. Moving forward, we aim to further examine the generalizability of our hypothesis across scales.

Acknowledgments

This work was funded by a Small Project Grant from the Centre for the Advanced Study of Collective Behaviour, the University of Konstanz (S20-06), by the University of Konstanz Committee on Research (FP031/19), and by the

Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC 2117–422037984. We thank Prof. Dr. Wolfgang Gaissmaier and Dr. Hansjörg Neth from the University of Konstanz, as well as Dr. Kenji Itao from the RIKEN Center for Brain Science, Japan, for their many helpful comments on this study.

References

Angus, S. D., & Newton, J. (2020). Collaboration leads to cooperation on sparse networks. *PLoS Computational Biology*, *16*(1), e1007557.

Barabâsi, A.-L., Jeong, H., Néda, Z., Ravasz, E., Schubert, A., & Vicsek, T. (2002). Evolution of the social network of scientific collaborations. *Physica A: Statistical mechanics and its applications*, *311*(3-4), 590–614.

Becker, J., Brackbill, D., & Centola, D. (2017). Network dynamics of social influence in the wisdom of crowds. *Proceedings of the national academy of sciences*, *114*(26), E5070–E5076.

Beekman, M., Sumpter, D. J., & Ratnieks, F. L. (2001). Phase transition between disordered and ordered foraging in pharaoh's ants. *Proceedings of the National Academy of Sciences*, *98*(17), 9703–9706.

Brackbill, D., & Centola, D. (2020). Impact of network structure on collective learning: An experimental study in a data science competition. *PLoS One*, *15*(9), e0237978.

Cantor, M., Chimento, M., Smeele, S. Q., He, P., Papageorgiou, D., Aplin, L. M., & Farine, D. R. (2021). Social network architecture and the tempo of cumulative cultural evolution. *Proceedings of the Royal Society B*, *288*(1946), 20203107.

Cantor, M., & Farine, D. R. (2018). Simple foraging rules in competitive environments can generate socially structured populations. *Ecology and evolution*, *8*(10), 4978–4991.

Collet, J., Morford, J., Lewin, P., Bonnet-Lebrun, A.-S., Sasaki, T., & Biro, D. (2023). Mechanisms of collective learning: How can animal groups improve collective performance when repeating a task? *Philosophical Transactions of the Royal Society B*, *378*(1874), 20220060.

Das, A., Gollapudi, S., & Munagala, K. (2014). Modeling opinion dynamics in social networks. *Proceedings of the 7th ACM international conference on Web search and data mining*, 403–412.

Dasaratha, K., & He, K. (2020). Network structure and naive sequential learning. *Theoretical Economics*, *15*(2), 415–444.

Deffner, D., Mezey, D., Kahl, B., Schakowski, A., Romanczuk, P., Wu, C. M., & Kurvers, R. H. (2024). Collective incentives reduce over-exploitation of social information in unconstrained human groups. *Nature Communications*, *15*(1), 2683.

Denrell, J. (2007). Adaptive learning and risk taking. *Psychological review*, *114*(1), 177.

Denrell, J. (2024). Adaptive sampling policies imply biased beliefs: A generalization of the hot stove effect. *arXiv preprint arXiv:2404.02591*.

Denrell, J., & March, J. G. (2001). Adaptation as information restriction: The hot stove effect. *Organization science*, *12*(5), 523–538.

Ghavasieh, A., & De Domenico, M. (2024). Diversity of information pathways drives sparsity in real-world networks. *Nature Physics*, *20*(3), 512–519.

Kao, A. B., & Couzin, I. D. (2014). Decision accuracy in complex environments is often maximized by small group sizes. *Proceedings of the Royal Society B: Biological Sciences*, *281*(1784), 20133305.

Langley, E. J., van Horik, J. O., Whiteside, M. A., & Madden, J. R. (2018). Individuals in larger groups are more successful on spatial discrimination tasks. *Animal Behaviour*, *142*, 87–93.

Lazer, D., & Friedman, A. (2007). The network structure of exploration and exploitation. *Administrative science quarterly*, *52*(4), 667–694.

Li, W., Zhang, X., & Hu, G. (2007). How scale-free networks and large-scale collective cooperation emerge in complex homogeneous social systems. *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics*, *76*(4), 045102.

Liden, R. C., Wayne, S. J., Jaworski, R. A., & Bennett, N. (2004). Social loafing: A field investigation. *Journal of management*, *30*(2), 285–304.

Lorenz, J., Rauhut, H., Schweitzer, F., & Helbing, D. (2011). How social influence can undermine the wisdom of crowd effect. *Proceedings of the national academy of sciences*, *108*(22), 9020–9025.

Mao, A., Mason, W., Suri, S., & Watts, D. J. (2016). An experimental study of team size and performance on a complex task. *PloS one*, *11*(4), e0153048.

Najar, A., Bonnet, E., Bahrami, B., & Palminteri, S. (2019). Imitation as a model-free process in human reinforcement learning. *bioRxiv*, 797407.

Rand, D. G., Nowak, M. A., Fowler, J. H., & Christakis, N. A. (2014). Static network structure can stabilize human cooperation. *Proceedings of the National Academy of Sciences*, *111*(48), 17093–17098.

Reina, A., Njougouo, T., Tuci, E., & Carletti, T. (2024). Speed-accuracy trade-offs in best-of-n collective decision making through heterogeneous mean-field modeling. *Physical Review E*, *109*(5), 054307.

Shirado, H. (2022). Individual and collective learning in groups facing danger. *Scientific Reports*, *12*(1), 6210.

Shirado, H., & Christakis, N. A. (2017). Locally noisy autonomous agents improve global human coordination in network experiments. *Nature*, *545*(7654), 370–374.

Shirado, H., & Christakis, N. A. (2020). Network engineering using autonomous agents increases cooperation in human groups. *Iscience*, *23*(9).

Smolla, M., & Akçay, E. (2019). Cultural selection shapes network structure. *Science advances*, *5*(8), eaaw0609.

Sumpter, D. J. (2006). The principles of collective animal behaviour. *Philosophical transactions of the royal society B: Biological Sciences*, *361*(1465), 5–22.

Toyokawa, W., & Gaissmaier, W. (2022). Conformist social learning leads to self-organised prevention against adverse bias in risky decision making. *eLife*, *11*.

Toyokawa, W., Whalen, A., & Laland, K. N. (2019). Social learning strategies regulate the wisdom and madness of interactive crowds. *Nature Human Behaviour*, 326637.

Watts, D. J., & Strogatz, S. H. (1998). Collective dynamics of 'small-world'networks. *nature*, *393*(6684), 440–442.