Social Learning Strategies Reconcile the Relationship Between Network Structure and Collective Problem Solving

Daniel Barkoczi
Mirta Galesic

SFI WORKING PAPER: 2015-12-052

Social learning strategies reconcile the relationship between network structure and collective problem solving

Daniel Barkoczi[1] and Mirta Galesic[2]

[1] Max Planck Institute for Human Development, Berlin, Germany
barkoczi@mpib-berlin.mpg.de
[2] Santa Fe Institute, New Mexico, USA
galesic@santafe.edu

Abstract. We study how different social learning strategies, composed of cognitively plausible rules that guide information search, stopping search and decision making, affect population-level performance in a collective problem-solving task. We show that different social learning strategies lead to remarkably different outcomes and demonstrate how these outcomes are affected by the communication networks agents are embedded in. We argue that understanding how communication networks affect collective performance requires taking into consideration the individual strategies used by agents. To illustrate this point we show how our findings can reconcile contradictory results in the literature on network structure and collective problem solving.

Keywords: social learning, NK model, exploration-exploitation, social networks

1 Introduction

The trade-off between exploration and exploitation lies at the heart of many problems faced by individuals, groups and organizations, who often need to decide whether to search for new, potentially better solutions (e.g., a technology, social institution, or a business strategy) or keep using an existing solution that works well [1, 2, 3, 4]. The right balance between exploration - searching for superior novel solutions - and exploitation - reaping benefits of existing solutions - is thought to be essential for adaptive behavior in humans and animals alike [5, 3, 4].

When individuals interact through social learning, (e.g., when solving problems collectively), this trade-off is manifested in the balance between innovation through individual learning from the environment and the imitation of existing solutions in the population [6, 7, 8, 9, 10]. Innovation (exploration) is essential both for tracking changes in the environment and for introducing novelty in the population, while imitation (exploitation) serves the purpose of diffusing good solutions in order to increase individual and group-level performance [11].

How do different social learning strategies affect behavior and performance in collective problem solving and how do these strategies interact with the social or organizational network in which learning takes place?

We address these two questions by modeling different social learning strategies as algorithms composed of three cognitively plausible building blocks, previously studied in the literature on adaptive individual cognition: rules that guide information search, stopping search and making a decision [12]. We study how agents using different strategies perform in simple and complex tasks, while embedded in social networks varying in structural properties that are known to affect the ease of information flow in communities [13, 14].

How network structure affects collective performance has been a long-standing interest of scholars of organizational behavior [15, 6, 16, 13, 14]. However, the answers that have emerged are somewhat contradictory. A number of studies have found that network structures that promote slower information diffusion (typically those that are less well connected) enhance collective performance in multi-peaked problem spaces, because they lead to higher levels of exploration and increase the chance of finding better solutions [6, 13, 17]. A recent study, focusing on the same problem, came to the opposite conclusion, finding that networks promoting faster information flow (typically those that are well connected) lead to better performance [14]. We argue that answering the question of how network structure affects performance requires studying how it interacts with the social learning strategies used by individuals.

We make two contributions. First, we demonstrate how the building blocks of social learning strategies can lead to strikingly different levels of performance. Second, we clarify and reconcile seemingly contradictory results in the literature by showing how social learning strategies and network structure interact to affect collective performance.

2 Modeling collective problem solving

We conceptualize collective problem solving as a task involving a group of agents repeatedly searching for solutions that improve individual and group-level performance. We follow several authors in modeling this problem as search on rugged landscapes [13, 14]. Each potential solution on the landscape consists of a number of components that are interdependent. Changing one component in a solution can affect the payoff-contribution of other components and, as a result, search on such landscapes is a form of combinatorial optimization. The main difficulty encountered when searching such landscapes is the presence of several local optima (solutions from which it is difficult or impossible to find better solutions). As a result, collectives face the challenge of finding good solutions without getting stuck on local optima.

Problem-solvers can engage in two primary activities: exploration through individual search or exploitation of known solutions via social learning. These two activities are complementary and require the right balance to achieve good performance.

Different network structures and social learning strategies can both affect the balance of exploration and exploitation in the population. Networks and social learning strategies that are fast at diffusing information both enable higher levels of exploitation at the expense of exploration. The opposite effect holds for networks and strategies that are slower at diffusing information. How different social learning strategies embedded in different network structures affect the balance of exploration and exploitation in the population is, therefore, crucial for understanding the factors determining collective performance.

2.1 Problem space

To design the problem space we use the NK model [18], which is a "tunably rugged" landscape denoted by N, the number of components that make up each solution, and K, the number of interdependencies between these components. N and K together determine the structure of the problem space where different solutions in the space have different payoffs. Consider a technology composed of different parts or an organization with different departmental configurations. Identifying a way to improve a technology or an organization's configuration depends both on its components and on the interdependencies between the components. For example, changing one component (e.g., increasing the number of departments in an organization) is likely to also have an impact on other components of the organization, and as a result, whether this change will increase overall performance depends on whether it also has a positive effect on the other components with which it interacts.

Depending on the number of interdependent components K, the landscape can be dominated either by a single global optimum ($K = 0$) in which case the payoffs of nearby solutions are highly correlated and local search is highly effective; by multiple local optima ($0 < K < N$), where payoffs of nearby solutions can have very different payoffs; or by almost completely uncorrelated landscapes where the payoffs obtained by local search become very similar to a random walk ($K = N - 1$).

To construct the environment we represent each solution in the environment by an N-length vector composed of binary strings, leading to a total of 2^N possible solutions in the problem space. The payoff of each solution is calculated as the average of the payoff contributions of each element. The payoff contribution of each element is a random number drawn from a uniform distribution between 0 and 1. In the case of $K = 0$, a simple average of the N elements is taken: $(1/N)\sum_{i=1}^{N} N_i$, whereas with $K > 1$, individual payoff contributions are determined by values of the $K-1$ other, interdependent elements, that is, $f(N_i \mid N_i, N_{i+1}, ..., N_K)$, where $f()$ is the payoff function and the total payoff is $(1/N)\sum_{i=1}^{N} f(N_i \mid N_i, N_{i+1}, ..., N_K)$. In other words, when $K = 0$, changing any single element of the solution will affect only the contribution of that element, whereas when $K > 0$, changing a single element will change the payoff contribution of the $K-1$ other elements. When $K = 0$, exploration of solutions through the modification of single components can prove effective, but as K increases, local exploration becomes less and less effective [5]. Figure 1 displays a simplified illustration of environments that vary in ruggedness. Panel A shows a simple environment where only one unique optimum exists and it is possible to reach this optimum by gradually modifying digits in one's solution. In contrast, Panel B shows a situation where several local optima exist, which means that agents can get stuck in a local optima and be unable to reach higher payoffs via local search.

We explore landscapes with $N = 15$ and $K = [0, 2, 4, 6, 7, 8, 10, 12, 14]$. Our choices for values of N and K are representative of the literature. Following several authors, we normalize the payoffs of different solutions by dividing them by the maximum obtainable payoff on a landscape $P_{Norm} = P_i/max(P)$ [13, 19]. The distribution of normalized payoffs tends to follow a normal distribution with decreasing variance as K increases. This implies that most solutions tend to cluster around very similar payoff values. Following [13] we use a monotonic transformation $(P_{Norm})^8$ to widen the distribution, making most solutions "mediocre" and only a few solutions "very good". This assumption does not change any of the results.

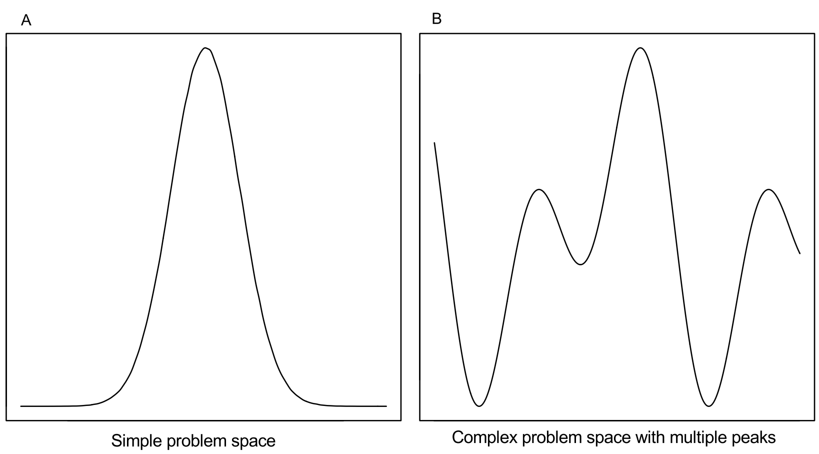

Fig. 1. Simplified illustration of the two environments studied. A: Simple environment with a single global optimum. **B:** Complex environment with multiple local optima and a global optimum. In the simple environment solutions one-digit apart from each other have very similar payoffs, so modifying single digits in a solution will eventually lead to the global optimum. In the complex environment payoffs of nearby solutions can be very different, so search by single digit modification can lead to local optima from which it is impossible to improve and, as a result, to find the global optimum.

We focus on two different problem spaces, a simple one with a single optimum ($N = 15$, $K = 0$) and a complex one with several local optima ($N = 15$, $K = 7$), but we have also tested several other values of K.

2.2 Social learning strategies

We focus on three frequently studied social learning strategies, namely: (1) *best member rule* [20], (2) *majority/plurality rule* (hereafter *conformity*) [21] and (3) *random copying rule* [10]. We formalize these strategies as algorithms composed of three basic cognitive building blocks: rules that guide information search, stopping search and making a decision [12];

(i) Search rule: all agents search randomly among the population of other agents.

(ii) Stopping rule: agents stop searching after looking up the solutions of s other individuals. We focused on two sample sizes: agents stop after collecting either a relatively small ($s = 3$) or a relatively large ($s = 9$) sample size.

(iii) Decision rule: agents either select the best performing agent (*best member*); select the most frequent solution (*conformity*; in case each solution is equally frequent, switch to individual learning); or select a random agent (*random copying*).

The combination of these rules produce five social learning strategies: *best member* with small and large samples, *conformity* with small and large samples, and *random copying*. Note that sample size does not affect the performance of the last strategy, so we study only one version that samples a single random individual.

As a benchmark, to see whether collective search improves performance over and above pure individual learning, we also consider a pure *individual learning* strategy, where agents explore possible solutions by randomly changing one digit in their current solution, but never copy other agents (see section on Simulation procedure for further details).

2.3 Network structure

We study a large number of networks that vary in several structural properties that have been proposed in the literature and have been found to affect collective problem solving. We consider a *fully connected network* where each agent is connected to every other agent in the population and eight additional network structures that were studied by [14]. Each of these eight networks maximize or minimize a specific loss function that corresponds to a specific network measure. The loss function that was minimized for each network is indicated by the network topology in Table 1. The obtained networks cover a broad spectrum of possible structures and the networks that were used in [13] and [14], the two studies that reached incompatible results. All networks have $N = 100$ nodes and a fixed degree of $d = 19$, except for the fully connected network where $d = N$. Table 1 shows the properties of each network. These different networks vary in several structural properties that have been previously shown to affect the speed of information flow. The top five networks are well connected and efficient at spreading information, while the bottom four networks are less well connected and hence less efficient at spreading information.

Table 1. Properties of the networks studied

Topology	Radius	Diameter	Closeness	Betweenness	Clustering	Constraint
Fully connected	1	1	1	1	1	1
Max max closeness	2	3	0.55	0.01	0.18	0.07
Min max closeness	2	3	0.55	0.01	0.18	0.07
Min mean betweenness	2	3	0.55	0.01	0.18	0.07
Min mean clustering	2	3	0.55	0.01	0.02	0.05
Max max betweenness	2	4	0.38	0.02	0.67	0.14
Max mean clustering	3	4	0.46	0.01	0.59	0.13
Max mean betweenness	4	7	0.35	0.02	0.56	0.13
Max var constraint	5	4	0.45	0.01	0.36	0.10

2.4 Simulation procedure

We simulated 100 agents. Agents searched the space of possible N-digit solutions in a given social network by modifying single digits in their current solutions in order to improve their performance. Agents by default engaged in social learning (exploitation) and switched to innovation (exploration) if the former did not prove successful. Specifically, on each time step agents went through the following steps:

(1) Implement social learning strategy composed of three building blocks:

(i) Search rule: search randomly among the population

(ii) Stopping rule: stop searching after looking up the solutions of s other individuals. In different simulations, all agents use one of two sample sizes: either a relatively small (s=3) or relatively large (s=9) sample size

(iii) Decision rule: in different simulations all agents either select the best performing agent (*best member*); select the most frequent solution (*conformity*)[3]; or select a random agent (*random copying*).

(2) Observe whether the solution identified via social learning produces a higher payoff than the current solution. If yes, switch to the alternative solution; otherwise go to Step 3.

(3) Engage in exploration by modifying a single digit in the current solution and observe whether it produces a higher payoff than the current solution. If yes, switch to the alternative solution; otherwise keep the current solution.

As a benchmark, in some simulations agents did not use social learning strategies but solely an *individual learning* strategy consisting of Step 3 above.

Steps (1)-(3) determine the complete strategy of an agent. It can be seen that the decision whether to explore through individual learning (step 3) or exploit through social learning (step 1) depends on whether social learning proves successful. As a result, reliance on exploration versus exploitation emerges naturally depending on the usefulness of social learning.

We iterate this procedure for t = 200 time steps and record the average payoff in the population on each time step separately for each combination of strategy,

[3] This implies selecting the majority/plurality solution in the sample. In case each solution is equally frequent, switch to individual learning.

network structure and problem space. Results reported are averaged across 1000 different draws of NK environments with the same parameters N and K.

3 Results

3.1 Performance of different social learning strategies

First we focus on the performance of different social learning strategies in a *fully connected* network in the absence of any structural properties. Figure 2 shows the average payoff achieved by each strategy over time for two different problem spaces: a simple one with a single optimum ($N = 15, K = 0$, left panel) and a complex one with several local optima and a global maxima ($N = 15, K = 7$, right panel). Table 2 summarizes results for all other values of K. Notice that the results are qualitatively the same for all values of $K > 0$.

Performance on simple problem spaces. In the simple problem space (where $K = 0$) all strategies eventually find the global optimum, however, strategies differ in the time it takes populations to converge (left panel of Figure 2). The *best member* strategies lead to the fastest convergence, followed by *conformity* and *random copying*. Since the simple problem space is dominated by only one optimum, the tension between exploration and exploitation is not so pronounced since eventually every individual will end up with the same solution.

Performance on complex problem spaces. The right panel of Figure 2 shows two striking results. First, the conformity strategy relying on small samples converges to the highest long-run outcomes, outperforming the *best member* strategies by a large margin. Second, the small-sample version of the *conformity* strategy outperforms the large-sample version.

In complex problem spaces (where $K > 0$), identifying good solutions is not straightforward since the space is dominated by several local optima. Here, *best member* strategies reach the highest short-run outcomes, but they quickly drive the whole population toward locally optimal solutions, from which point individual exploration is no longer able to find better solutions. As a result the whole population gets stuck in an inferior state. Compared to *best member*, the *conformity* strategies converge more slowly and, therefore, lead to higher levels of exploration in the population.

While the large-sample version of the *conformity* strategy performs poorly, the small-sample version converges to the highest long-run outcomes, outperforming the *best member* strategies by a large margin. The underlying reason is that the *conformity* strategy is less efficient at finding a good solution, leading to more exploration of the problem space. In addition, the small-sample version of *conformity* is able to diffuse good solutions that are discovered later-on, because its relatively large sampling error makes it more likely that good solutions identified by only a few individuals in the population will appear as the most frequent solution in an individual sample. This allows infrequent but superior solutions

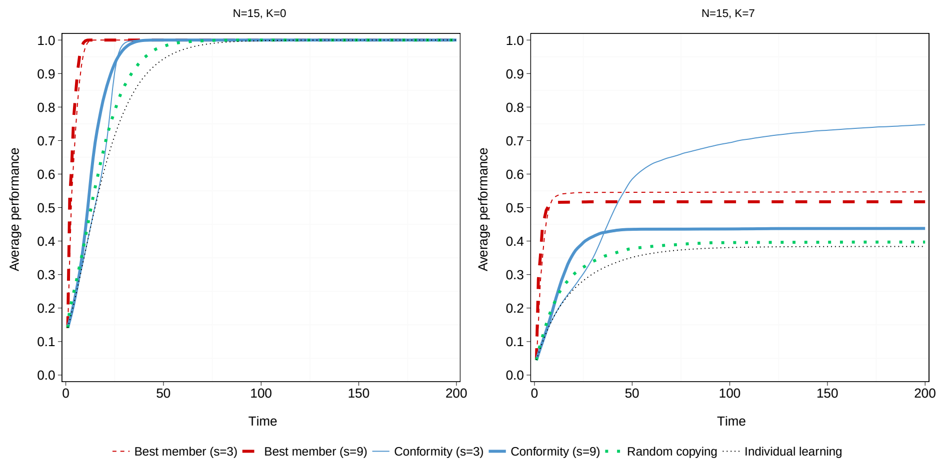

Fig. 2. Performance over time for different strategies in a fully connected network. N denotes the number of components of the system and K represents the number of interactions between the components; s stands for sample size. Left: simple environment with a global optimum, Right: complex environment with multiple local optima.

discovered later on to diffuse through the population (see also [22]). In the long run the small-sample version of *conformity* is able to reach the highest outcomes, as the group is able to search extensively as well as to converge on high fitness solutions over time. In contrast, the smaller sampling error of the large-sample version of *conformity* makes this strategy unable to diffuse infrequent but useful information and leads to poor performance.

That the superiority of the small-sample *conformity* strategy does not stem purely from the fact that it makes individuals explore individually when social learning cannot identify good solution can be seen from the fact that it outperforms two simple benchmarks, namely pure *individual learning* and *random copying*. This pattern of results was replicated in all complex landscapes ($N > K > 0$, see Table 2).

Random copying, like *best member*, engages in high levels of exploitation in the beginning and drives agents to several locally optimal solutions. However, since this decision rule is not biased towards any criteria related to success (e.g. best option, most frequent option) it is not able to drive the population to good solutions.

Finally, all social learning strategies outperformed pure *individual learning*, replicating the finding that collectives outperform individuals in this task [13, 14, 23].

Taken together, these results indicate that different social learning strategies lead to different patterns of explorative and exploitative behavior over time. Strategies such as the *best member* rule lead to high levels of exploitation and drive the population toward local optima. Other strategies such as the *conformity* rule promote higher levels of exploration and enable the population to find higher-payoff solutions. The extent to which different strategies prove useful also depend crucially on their building blocks (search, stopping and decision rules).

Since different strategies can achieve remarkably different performance within the same network, it is important to study how this difference plays out and interacts with with other network structures that are also known to affect performance through exploration and exploitation [13, 14].

3.2 Interaction of network structure and decision strategy

We now proceed to the question of how network structure and decision strategies interact. We investigate how strategies perform in different networks, and what determines the superiority of well or less well connected networks. We verify our conclusions above in more realistic network structures, and reconcile two contradictory results in the literature. Lazer and Friedman [13] use an agent-based simulation to compare a fully connected network to a locally connected lattice and found that the locally connected network outperformed the fully connected network in the long run. This result implies that networks that lead to slower information spread lead to better outcomes. In contrast, Mason and Watts [14] report a behavioral experiment with eight different networks (see Table 1). They find that networks that are faster at spreading information in the population outperform slower networks. What is driving this difference in results? Given the computational nature of the study by Lazer and Friedman [13] we know the exact strategy their agents relied on: the *best member* rule with a sample size of two. However, given the data available we cannot tell what strategies participants used in the study of Mason & Watts [14].

Here we show that both results can be obtained depending on the social learning strategies that agents use in a given network. To demonstrate this point we re-run the study reported in the previous section in all networks reported in Table 1. We focus on the two best performing social learning strategies, the *best member* and *conformity* strategies with small samples.

We focus on the eight different networks studied by [14]. As in our previous simulations, we set $N = 100$ and $d = 19$ (these networks have the same degree to node ratio as in the study of [14]).

Figure 3 shows the average payoff achieved by groups in different networks. The left panel shows the average performance of the eight networks when agents rely on the *best member* rule, while the right panel shows the same performance when agents rely on the *conformity* rule. Results in the left panel replicate the findings of [13] who find that less well connected networks outperform better connected ones, while results in the right panel replicate the findings of [14] who report the opposite result. By comparing the two panels we can also see that the *conformity* strategy outperforms the *best member* strategy in each network.

Our findings demonstrate that both types of networks can lead to superior performance depending on the social learning strategies used by individuals in the group. The underlying explanation is the following. Network structure and social learning strategies both affect the levels of exploration and exploitation in the population. If both strategy and network promote high levels each of the two behaviors, performance is likely to drop, however, if network and strategy promote opposite effects, performance is likely to rise.

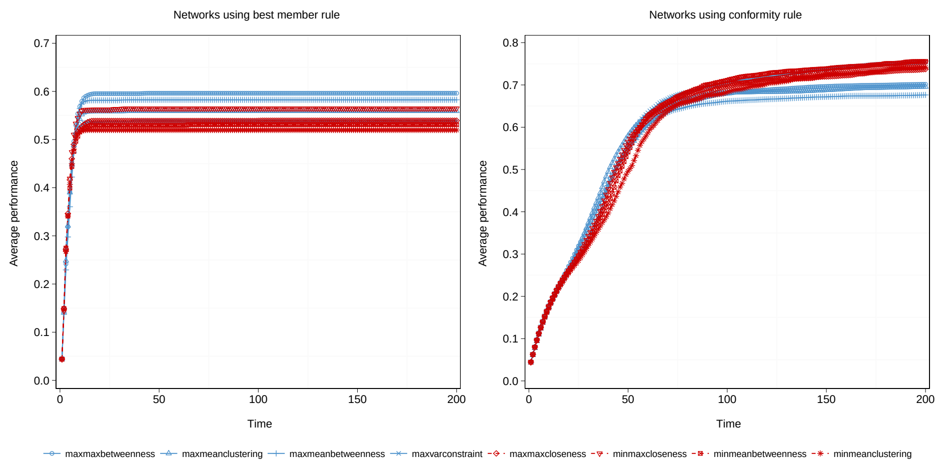

Fig. 3. Collective performance on networks using when agents use the best member rule (left panel) and the conformity rule (right panel). Well connected networks are marked in red and less well connected networks in blue. The left panel replicates the findings from Lazer & Friedman (2007)[13] while the right panel replicates the findings of Mason & Watts (2012)[14]

3.3 Sensitivity Analyses

We perform two sensitivity checks. First, we study how our main results are affected by environmental turbulence and second, we study the best sample size for the *best member* and *conformity strategies*. For both checks we use the fully connected network.

Changing task environments. We model two types of environmental change. In the first case we completely regenerate the NK landscape half-way through the simulation, forcing agents to re-learn everything that was adaptive in the past. This represents an environment that changes rarely but drastically (for an example see [24]). In the second case we redraw the fitness contribution of a randomly selected single digit in the solution space every 40th time step. This represents a more frequent, but less drastic change. From Figures 4 we conclude that environmental change does not alter our main conclusions.

Best sample size. We focus on identifying the best sample size for the *best member* and *conformity* strategies, sample size has no effect on *random copying*.

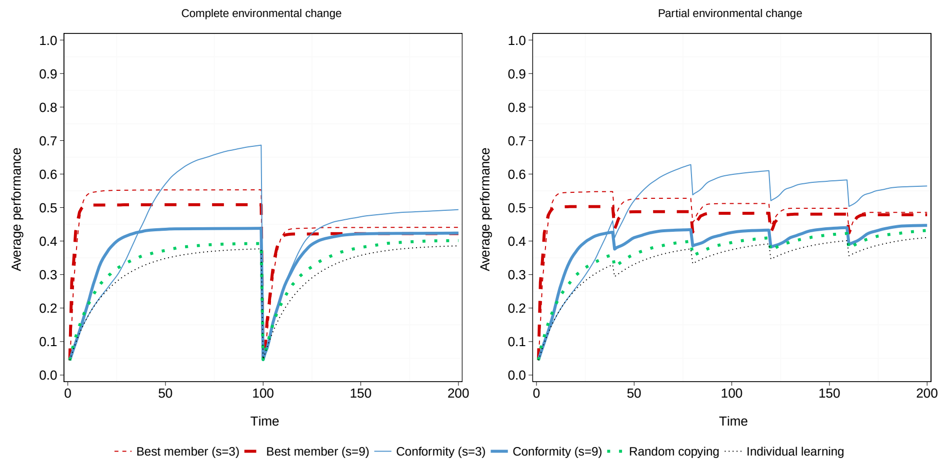

Fig. 4. Performance over time for different strategies in changing environments.Left panel: Rare but drastic environmental change. Right panel: Frequent but less drastic environmental change.

For the *best member* strategy the best sample size turns out to be $s = 2$, however, the difference between different sample sizes is relatively small[4]. Therefore we choose to keep sample size of $s = 3$ in the main text to make it directly comparable to the *conformity* strategy that also uses a sample size of $s = 3$. For *conformity* the best sample size is $s = 3$.

4 Discussion

We studied how social learning strategies affect collective performance in different social network structures and under varying levels of task complexity.

We modeled strategies as algorithms composed of three cognitively plausible building blocks - that is, rules that guide search, stopping search, and making a decision - and studied how these rules affect the performance of different strategies.

We asked two questions. First, how do different social learning strategies affect behavior and performance in collective problem solving? We found that *best member* strategies reach the best performance in the short run, but a small amount of *conformity* (achieved by relying on small samples) ensures the highest long-run outcomes whenever task environments are complex. The intuition underlying these results is the following. The *best member* strategies are fast at diffusing useful information and, therefore, quickly drive the population toward locally optimal solutions. Reliance on small-samples performs slightly better than

[4] Note that the best member strategy with a sample size of 1 would correspond to the *random copying strategy*.

large-samples, because it leads to slightly slower convergence and thereby allows the population to explore and find optima that have higher payoffs.

Second, how do these strategies interact with the network structure in which learning takes place? Our results indicate that networks promoting faster information diffusion outperform the slower networks when agents use *conformity* strategy. However, opposite is the case when agents use *best member* strategy: here, slower networks outperform faster ones. This shows that collective performance depends both on the network structure agents are embedded in as well as the social learning strategies they use. These findings enabled us to clarify and reconcile seemingly contradictory findings from the literature, by showing that both well connected and less well connected networks can be beneficial for the same task, depending on the social learning strategies used by individuals [6, 13, 17, 14].

Our study has broad implications for organizational learning, technological innovation and the diffusion of innovations. Most studies of exploration and exploitation in organizations focus on how to design the external environment to make firms more adaptive [25, 6, 13, 19]. Our study highlights that it is also important to consider the individual strategies used by agents and organizations. In addition our study shows that interventions aimed at changing the social environment without paying attention to the individual-level strategies might not produce the desired effect.

Research on technological innovation has highlighted the combinatorial nature of innovation with most new inventions being recombinations of existing innovations [26, 27]. Much of this research has focused on how innovation occurs, whereas there has been very little attention devoted to the co-evolution of innovation and imitation. Our study identifies situations where imitation can both help and hinder the development of technological innovation.

Several open questions remain to be addressed. In line with previous studies we focused on the NK landscape as a form of a tunably rugged landscape. The extent to which our results (and other results from the literature) would apply to other landscape problems is a question for future research. We also assumed for the sake of clarity that populations rely on a single social learning strategy. Future research should address the dynamics of exploration and exploitation in a population using multiple strategies at the same time. Our model could also be tested empirically. There are only a handful of studies on how people behave in combinatorial optimization problems that have a rugged structure and we know very little about how these results translate to other problems [28, 29, 30, 31].

Acknowledgments

We thank John Miller, Scott Page and participants of the 21st Graduate Workshop in Computational Social Science Modeling and Complexity Workshop held at the Santa Fe Institute for helpful comments and the Max Planck Institute for Human Development for financial support.

Table 2. Strategy performance at the final time step ($t = 200$) in different environments.

Strategy		Environment						
		$N = 15$, $K = 2$	$N = 15$, $K = 4$	$N = 15$, $K = 6$	$N = 15$, $K = 8$	$N = 15$, $K = 10$	$N = 15$, $K = 12$	$N = 15$, $K = 14$
Best member ($s = 3$)	mean	0.8522	0.705	0.600	0.513	0.444	0.408	0.376
	min	0.288	0.179	0.171	0.150	0.131	0.081	0.111
	max	1	1	1	1	1	1	1
	st.dev	0.166	0.220	0.218	0.200	0.174	0.163	0.148
Best member ($s = 9$)	mean	0.823	0.688	0.567	0.479	0.424	0.381	0.369
	min	0.249	0.187	0.118	0.094	0.116	0.102	0.109
	max	1	1	1	1	1	1	1
	st.dev	0.185	0.233	0.220	0.197	0.181	0.154	0.151
Conformity ($s = 3$)	mean	0.978	0.915	0.807	0.681	0.571	0.461	0.377
	min	0.370	0.446	0.205	0.254	0.219	0.145	0.119
	max	1	1	1	1	1	1	1
	st.dev	0.062	0.125	0.176	0.188	0.185	0.168	0.147
Conformity ($s = 9$)	mean	0.714	0.563	0.469	0.491	0.356	0.314	0.280
	min	0.222	0.230	0.208	0.197	0.166	0.146	0.126
	max	1	1	1	1	1	1	1
	st.dev	0.148	0.132	0.112	0.092	0.103	0.082	0.072
Random copying	mean	0.681	0.517	0.428	0.369	0.314	0.274	0.235
	min	0.090	0.102	0.093	0.051	0.052	0.044	0.061
	max	1	1	1	1	1	1	1
	st.dev	0.236	0.237	0.199	0.189	0.166	0.140	0.115
Individual learning	mean	0.681	0.518	0.420	0.354	0.302	0.261	0.228
	min	0.205	0.285	0.156	0.144	0.152	0.127	0.114
	max	0.954	0.718	0.638	0.515	0.486	0.418	0.336
	st.dev	0.092	0.074	0.066	0.058	0.051	0.044	0.038

Bibliography

[1] Gupta, A.K., Smith, K.G., Shalley, C.E.: The interplay between exploration and exploitation. Academy of Management Journal **49**(4) (2006) 693–706

[2] Holland, J.H.: Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence. University of Michigan Press (1975)

[3] March, J.G.: Exploration and exploitation in organizational learning. Organization Science **2**(1) (1991) 71–87

[4] Mehlhorn, K., Newell, B.R., Todd, P.M., Lee, M., Morgan, K., Braithwaite, V.A., Hausmann, D., Fiedler, K., Gonzalez, C.: Unpacking the exploration-exploitation tradeoff: A synthesis of human and animal literatures. Decision (2015)

[5] Levinthal, D.A.: Adaptation on rugged landscapes. Management Science **43**(7) (1997) 934–950

[6] Fang, C., Lee, J., Schilling, M.A.: Balancing exploration and exploitation through structural design: The isolation of subgroups and organizational learning. Organization Science **21**(3) (2010) 625–642

[7] Kameda, T., Nakanishi, D.: Cost–benefit analysis of social/cultural learning in a nonstationary uncertain environment: An evolutionary simulation and an experiment with human subjects. Evolution and Human Behavior **23**(5) (2002) 373–393

[8] Kameda, T., Nakanishi, D.: Does social/cultural learning increase human adaptability?: Rogers's question revisited. Evolution and Human Behavior **24**(4) (2003) 242–260

[9] Rendell, L., Boyd, R., Cownden, D., Enquist, M., Eriksson, K., Feldman, M.W., Fogarty, L., Ghirlanda, S., Lillicrap, T., Laland, K.N.: Why copy others? insights from the social learning strategies tournament. Science **328**(5975) (2010) 208–213

[10] Rogers, A.R.: Does biology constrain culture? American Anthropologist **90**(4) (1988) 819–831

[11] Boyd, R., Richerson, P.J.: Culture and the evolutionary process. University of Chicago, Chicago (1985)

[12] Gigerenzer, G., Todd, P.M., the ABC Research group: Simple heuristics that make us smart. Oxford University Press (1999)

[13] Lazer, D., Friedman, A.: The network structure of exploration and exploitation. Administrative Science Quarterly **52**(4) (2007) 667–694

[14] Mason, W.A., Watts, D.J.: Collaborative learning in networks. Proceedings of the National Academy of Sciences **109**(3) (2012) 764–769

[15] Bavelas, A.: Communication patterns in task-oriented groups. Journal of the acoustical society of America (1950)

[16] Guetzkow, H., Simon, H.A.: The impact of certain communication nets upon organization and performance in task-oriented groups. Management science **1**(3-4) (1955) 233–250

[17] Mason, W.A., Jones, A., Goldstone, R.L.: Propagation of innovations in networked groups. Journal of Experimental Psychology: General **137**(3) (2008) 422

[18] Kauffman, S.A., Levin, S.: Towards a general theory of adaptive walks on rugged landscapes. Journal of Theoretical Biology **128**(1) (1987) 11–45

[19] Siggelkow, N., Rivkin, J.W.: Speed and search: Designing organizations for turbulence and complexity. Organization Science **16**(2) (2005) 101–122

[20] Rivkin, J.W.: Imitation of complex strategies. Management Science **46**(6) (2000) 824–844

[21] Henrich, J., Boyd, R.: The evolution of conformist transmission and the emergence of between-group differences. Evolution and Human Behavior **19**(4) (1998) 215–241

[22] Barkoczi, D., Galesic, M.: A sampling approach to conformist social learning. Available at SSRN 2479414 (2015)

[23] Shore, J., Bernstein, E., Lazer, D.: Facts and figuring: An experimental investigation of network structure and performance in information and solution spaces. Organization Science (2015)

[24] Whitehead, H., Richerson, P.J.: The evolution of conformist social learning can cause population collapse in realistically variable environments. Evolution and Human Behavior **30**(4) (2009) 261–273

[25] Csaszar, F.A., Siggelkow, N.: How much to copy? determinants of effective imitation breadth. Organization Science **21**(3) (2010) 661–676

[26] Arthur, W.B.: The nature of technology: What it is and how it evolves. Simon and Schuster (2009)

[27] Solée, R.V., Valverde, S., Casals, M.R., Kauffman, S.A., Farmer, D., Eldredge, N.: The evolutionary ecology of technological innovations. Complexity **18**(4) (2013) 15–27

[28] Billinger, S., Stieglitz, N., Schumacher, T.R.: Search on rugged landscapes: An experimental study. Organization Science **25**(1) (2013) 93–108

[29] Garcia-Retamero, R., Takezawa, M., Gigerenzer, G.: How to learn good cue orders: When social learning benefits simple heuristics. In: Proceedings of the 28th annual conference of the cognitive science society. (2006) 1352–1358

[30] Mesoudi, A.: An experimental simulation of the "copy-successful-individuals" cultural learning strategy: adaptive landscapes, producer–scrounger dynamics, and informational access costs. Evolution and Human Behavior **29**(5) (2008) 350–363

[31] Wisdom, T.N., Song, X., Goldstone, R.L.: Social learning strategies in networked groups. Cognitive Science **37**(8) (2013) 1383–1425