

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Group problem solving: Diversity versus diffusion

Permalink

<https://escholarship.org/uc/item/0zq5j37r>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 46(0)

Authors

Jonard, Nicolas

Reijula, Samuli

Marengo, Luigi

Publication Date

2024

Peer reviewed

Group problem solving: Diversity versus diffusion

Nicolas Jonard (nicolas.jonard@uni.lu)

Department of Economics and Management
University of Luxembourg
Luxembourg, L1359, Luxembourg

Samuli Reijula (samuli.reijula@helsinki.fi)

Theoretical Philosophy, PO Box 24
00014 University of Helsinki, Finland

Luigi Marengo (lmarengo@luiss.it)

Department of Business and Management
LUISS University
Rome, 00197, Italy

Abstract

Several recent contributions to the research on group problem solving suggest that reducing the connectivity between agents in a social network may be epistemically beneficial. This notion stems from the idea that collective problem-solving behavior may benefit from the transient diversity in agents' beliefs due to increased individual exploration and decreased social influence. At the same time, however, lower connectivity hinders the diffusion of good solutions between network members. Our simulation findings shed light on this trade-off. We identify conditions under which the less-is-more effect is likely to manifest. Our findings suggest that a community consisting of semi-isolated groups could provide an answer to the tension between diversity and diffusion.

Keywords: group problem solving; cognitive diversity; structural isolation; agent-based simulation

Introduction

Several recent contributions to research on group problem solving suggest that less may be more when it comes to communication between problem-solving agents (Lazer & Friedman, 2007; Zollman, 2010; Fang, Lee, & Schilling, 2010). In other words, less information sharing between the members of a group or a collective may lead to better problem-solving outcomes.

The less-is-more effect is relevant to the practical question of how to best structure the interaction between the components of a distributed cognitive system, i.e., the members of a team or a broader research community. Skunkworks divisions at technology companies are an example of cutting network ties in order to promote the quality of problem-solving activities: structurally isolating a R&D department from the rest of the organization has been used as a means to preserve its divergent but valuable ideas (Fang et al., 2010).

Zollman (2010) explored a similar idea in the context of scientific research, and suggested that reducing the connectivity between agents in the social network of the scientific community may be epistemically beneficial.¹ Zollman applied his model to historical episodes of theory choice. The model suggests that introducing *transient diversity* of beliefs, either by reducing connectivity or by equipping agents with strong

initial beliefs, increases the probability that true consensus can ultimately be reached.²

But how little connectivity is too little – can it really be the case that the less problem-solving agents communicate, the better the aggregate outcome? The literature suggests that two opposing forces are at play: On the one hand, the more strongly connected an epistemic network, the more likely it is that beneficial diversity washes out from the population. On the other hand, in a very fragmented network, heterodox but valuable beliefs will not spread and hence will not influence the consensus view (Fang et al., 2010). In both cases, transient diversity fails to promote good epistemic outcomes.

In this paper we address this trade-off by building a series of computational experiments. Compared to earlier work on the topic, our findings suggest a more complicated picture of the less-is-more effect. First, we fail to detect the effect in large groups. Secondly, in semi-isolated groups, we find a non-monotonic effect: Although under certain conditions the collective performance worsens as density increases (as suggested by the less-is-more effect), under different conditions more seems to be more, i.e., increases in density improve collective performance. From these findings, we conclude that the preconditions and dynamics of the less-is-more effect are not yet fully understood.

Our findings suggest, however, that splitting a community of problem solvers into semi-isolated clusters helps to strike a balance fostering epistemically beneficial transient diversity.

Existing literature and theoretical background

March's seminal *Exploration and exploitation in organizational learning* provided reasons to think that the less-is-more effect may be explained by the exploration-exploitation trade-off (March, 1991; Mehlhorn et al., 2015). Also more recent contributions to the literature suggest that reducing connectivity between agents improves problem-solving outcomes by promoting exploration in the community and thus preventing premature convergence to sub-optimal outcomes.

²For an overview of social and cognitive diversity in science, see Rolin, Koskinen, Kuorikoski, and Reijula (2023).

¹See also Kummerfeld and Zollman (2016).

Lazer and Friedman (2007) combined March’s general approach with an NK landscape. They found that when dealing with a complex problem, more efficient networks perform well in the short run but worse in the long run. Lazer & Friedman’s findings suggest that an inefficient network maintains diverse beliefs in the system and thus supports exploration.

Zollman (2010) reached similar conclusions in a setup where a community of agents who are able to learn from their network neighbors faced a bandit problem. According to Zollman’s model, the probability of correct convergence to the better bandit arm on dense graphs is lower than on sparse ones.

Using a variation of March’s (1991) model, Fang et al. (2010) examined whether a beneficial exploration-exploitation balance could be struck by relying on social network structures consisting of semi-isolated cliques. Findings by Barkoczi and Galesic (2016) suggest a more complicated picture where the effect of network connectivity is moderated by the social learning rule used by agents.

In addition to model-based findings, there is some empirical work on the topic. The empirical evidence for the less-is-more effect is mixed: While Mason, Jones, and Goldstone (2008) find some support, the effect seems to be sensitive to various task properties. Similarly, the findings by Shore, Bernstein, and Lazer (2015) confirm an ambiguous effect of connectivity: While it can encourage network members to generate more non-redundant information, it may also hinder exploration (see also Yahosseini, Reijula, Molleman, & Moussaid, 2018). Mason and Watts (2012)’s findings are similarly inconclusive.

The research question

In this article, we address the less-is-more puzzle by combining insights from Zollman (2010) and Fang et al. (2010). Like Zollman, we study social learning by (myopic) Bayesian agents facing a two-armed bandit task. However, in our replications, we depart from Zollman in studying also larger networks, and by varying network size and density independently, so as to tease apart the influence of the two properties.

After this first step, conducted in the spirit of replication, we consider a new setting in which we vary network density and architecture independently, to explore the value of problem-solving groups.

Imagine a large community consisting of several interconnected “cliques.” The underlying structure could represent a technological or scientific community, within which a clique is a group of inventors working together on a topic, jointly patenting their findings or coauthoring scientific papers. Alternatively, the underlying structure could represent an organization, within which cliques are departments or teams jointly trying to make sense of their environment and improve their performance. In either case, the structure consists of sparsely interconnected groups whose members, inside a group, are densely connected.³

³Although formally, in graph theory, a clique is a subset of nodes such that every two distinct nodes in the clique are connected, we use

Inspired by Fang and colleagues’ findings, we hypothesize that dividing the population of agents into such cliques and manipulating the proportion of links held by an agent within her clique relative to the those held outside her clique could be a way to find the sweet spot between exploration and exploitation, transient diversity and efficient diffusion. While dense cliques will accumulate knowledge and converge rapidly on one alternative (which may not be the superior one), inter-clique connectivity brings together the perspectives of the different groups and can correct sub-optimal decisions that have been made in some of them.

To study semi-isolated groups of inquirers, we examine agents organized as a connected caveman graph (Watts, 1999). The caveman graph provides a tunable algorithm that allows for interpolation between two extremes: isolated groups of maximally connected agents on one end, and random graphs on the other, by simply controlling the proportion p_r of edges in the initial structure that are randomly rewired.

The model

In order to promote transparency and readability, the description of our model follows the *Overview, Design concepts and Details (ODD) protocol* for describing individual- and agent-based models (Grimm et al., 2020).

Purpose This model is designed to better delineate the conditions under which the less-is-more effect is at work in groups of different sizes and levels of connectivity, and to provide an extension to a richer social structure of interconnected groups.

Entities and state variables *Two-armed bandit.* The problem-solving population is faced with the task of choosing the better arm of a two-armed bandit. Each arm corresponds to a Bernoulli bandit $j \in \{1, 2\}$, which pays 1 with probability p_j and 0 with complementary probability $1 - p_j$.

Agents. The problem-solving agents explore the choice alternatives by sampling the two arms. Social learning takes the following form: at the time of choosing an arm, an agent pools the information collected through her own sampling and the information obtained by her network neighbors’ sampling. Based on her own and her neighbors’ outcomes, the agent updates her beliefs following Bayes’ rule. Each agent utilises a greedy decision rule, and pulls the arm which yields the highest expected payoff given the current beliefs (see Zollman, 2010).

Network. The informational connections between the agents are represented as edges in an undirected network. Community size and network density are varied independently. In our first simulation study, community size $n \in \{3, 6, 8, 12, 24, 48, 96\}$.⁴ The network density is varied from its minimum possible value to the maximal value of 1.⁵ In our second simulation

the term more loosely, and will consider cliques of various densities.

⁴Zollman (2010) goes up to $n = 11$, and his explorations on density are for small networks of size $n = 6$.

⁵There is a minimum possible value to density in a network of size n when the network must be connected: connectivity requires at least $n - 1$ links, and there are $n(n - 1)/2$ possible links, so density is at

study, semi-isolated groups are modeled as connected caveman graphs (Watts, 1999). We use $n = 96$ and a cave size n_c in $\{3, 4, 6, 8, 12, 16, 24, 32, 48, 96\}$. If cave size is n_c , then n/n_c caves can be formed.

On the population level, the most important state variables tracked are the proportion of correct learners in each trial, and the time to convergence.

Process overview and scheduling Each time period, each agent pulls one of the two arms, and then by using Bayes' rule, agents pool information from their neighbors (augmented with their own information) and update their belief distribution. This process simply iterates. Through the process, the community of inquirers eventually tends to form a consensus. Each run of the simulation lasts for 10,000 time periods, and we present results after 2,000 and 10,000 periods.

Design concepts The first *core design principle* of the model derives from the bandit paradigm (Berry & Fristedt, 1985; Sutton & Barto, 2018), combined with social learning in networks (Bala & Goyal, 1998; Zollman, 2010). The second core idea is semi-isolation in networks (Fang et al., 2010) implemented as connected caveman graphs (Watts, 1999). *Adaptation* in agent behavior is modeled as Bayesian learning. The *objective* of each individual agent is to discover the best arm of the two-armed bandit. The agents can *sense* their environment in two ways, by pulling an arm of the machine, and by receiving information about outcomes of other agents' pulls from their network neighbors (= *interaction*). On the *collective* level, the most important properties to track are the proportion agents converged to the better arm, and the time to convergence.

Initialization Agents' prior beliefs are beta-distributed, with random initial beliefs characterized by agent-specific parameters $a_{i,j}$ and $b_{i,j}$ uniformly and independently distributed over $[0, 4]$, as in Zollman (2010). The true success probabilities for the arms of the bandit are $\pi_1 = .5$ for the first arm, and $\pi_2 = .49$ for the second arm.

In each of the 1,000 independent replications (simulation runs) that we consider, a random network is generated with the appropriate size and density in the first study, and with the appropriate number of caves, cave density and amount of rewiring in the second study. Each node of the network then represents an agent.

least equal to $(n-1)/[(n(n-1)/2)] = 2/n$. To illustrate, in a network of size $n = 4$, there can be at most $4 \times 3/2 = 6$ edges (the complete network obtains when the first agent forms a pair with each of the 3 others, the second agent forms a pair with the 2 remaining others, and the third agent forms a pair with the last remaining agent she is not yet connected to). The network requires also at least 3 edges to be connected. Therefore, the density of any connected network with 4 nodes belong to the set $\{1/2, 2/3, 5/6, 1\}$.

Input data, and submodels The simulations do not employ empirical data as input. The relevant (sub)model structure has been described above.

Simulations and results

Simulation study 1

In our first simulation study, we replicate findings from Zollman (2010), and extend them to larger groups. The main finding is that in our setup, we do not observe the less-is-more effect in large communities when we consider a 2,000-period learning horizon.




Figure 1: The accuracy of learning in communities of various sizes and densities ; each curve is for one community size $n = 3, 6, 8, 12, 24, 48, 96$ over a learning horizon of 2,000 periods.

In Figure 1, the less-is-more effect is present for communities of size less than or equal to 8, but for communities of size 12 and larger, there is a clear non-monotonic relationship between cognitive performance and density, which peaks when density is approximately 35%.

In general, diversity persists longer in less dense communities, fostering short run exploration but hindering long run exploitation, thus degrading aggregate performance. While this is not an issue in small communities, in large communities diversity persists in excess amounts when density is too low, which prevents overall diffusion of the best alternative — because of this tension, an ‘optimal’ amount of diversity obtains for intermediate values of density.⁶

To explore further this tension between diversity and diffusion, we present in Figure 2 the relationship between cognitive

⁶For all of the combinations of parameter values examined, simulations were repeated 1,000 times (yielding 1,000 independent replications), and we present average values.

performance and density over a 10,000-period learning horizon.




Figure 2: The accuracy of learning in communities of various sizes and densities; each curve is for one community size $n = 3, 6, 8, 12, 24, 48, 96$ over a learning horizon of 10,000 periods.

With more learning and interaction opportunities, an alternative is identified as the superior one more reliably and diffuses more widely in the community. The less-is-more effect is now present for communities of size up to 12, though a non-monotonic relationship between cognitive performance and density remains visible for larger community sizes. Interestingly, the ‘optimal’ amount of diversity (in the sense of supporting both enough exploration to identify the superior alternative and enough later convergence) is now obtained for smaller (though not minimal) values of density (closer to 20%, relative to the 35% that seemed ideal when only 2,000 periods were considered). Not only has the peak shifted to the left (when a peak exists), but it is worth noting that (unsurprisingly) all communities display better average performance when given more learning time. It is likely that increasing further the learning horizon would contribute to strengthening the less-is-more effect and it would improve the relative performance of sparse structures. However, prolonged exploration will have diminishing returns: smaller and smaller performance gains will obtain at the expense of ever-increasing learning efforts. Whether the marginal benefits of prolonging exploration exceed the associated cost is thus an open question.

While Figures 1 and 2 represent average behaviour, it is interesting to examine in more detail how the entire distribution of aggregate performance changes when density changes,

in the specific case $n = 48$. For each of the values of density that are considered, we look at the corresponding set of 1,000 replications and tabulate (across the set of replications) the proportion of successful learners after 2,000 periods. The resulting two-dimensional histogram is represented as a surface plot over the density-accuracy plane. Picking any level of density, and then moving parallel to the accuracy axis considering any value between 0 and 1 gives the count of replications (over the 1,000 replications that were run) that have produced the corresponding value of accuracy. In that sense, the surface represents a family of count histograms parameterized by the value of density.




Figure 3: The accuracy of learning in communities of size $n = 48$ for different densities: at each density, the distribution of outcomes in period 2,000 over the 1,000 replications is depicted.

We observe a transition from a unimodal distribution of outcomes peaking near 0.8 for low values of density, to a bimodal distribution of outcomes, with two marked peaks at 0 and 1 and increasing mass on the lower mode of the distribution (corresponding to all agents picking the wrong arm) when density approaches its maximal value of 1. This emphasizes how two seemingly similar macro-level outcomes can represent very different micro-level situations. Considering the curve associated with $n = 48$ in Figure 1, the average accuracy obtained for the lowest density value of $2/48 = 0.042$ and the average accuracy for a density value nearing 80% (the exact value is $877/[48(48 - 1)/2] = 0.78$) correspond to the same level of average performance for the system (close to 80%). However, in the former case, in any simulation run many agents do well (the mass of the distribution is fairly concentrated around .8, so the proportion of correct learners

is often high, and very rarely low), whereas in the latter case, though in roughly 3 out of 4 replications the superior arm is identified by the entire community it is also the case that in every 4th replication, the entire community selects the wrong arm. Though the mean is (roughly) preserved, variance in replication outcomes starkly differs depending on the density configuration that is considered.

In terms of governance of the system and policy decisions more generally, the second situation we described can be seen as much riskier than the first one: while in the first situation an important part (more than half) of the whole community always ‘has it right’, in the second situation all ‘have it wrong’ with probability 25%. This vividly illustrates the tension between diversity and diffusion.

Simulation study 2

In our second simulation, we examine problem-solving performance in a caveman graph consisting of dense semi-isolated groups whose ties are rewired randomly with a probability we control.




Figure 4: Three configurations for a minimally connected caveman graph of $n = 30$ agents distributed over 5 caves of size $n_c = 6$, with increasing rewiring probability (left to right, $p_r = 0$, $p_r = 0.1$ and $p_r = 1$) showing the transition from a locally highly clustered structure to a uniform random graph.

Figure 4 illustrates how network structure changes when the rewiring probability p_r is varied. In this experiment, our tuning parameter is p_r , the probability of randomly rewiring an edge. Starting from an initial caveman graph consisting of complete caves of size n_c minimally connected by clique-spanning ties, we consider each edge sequentially and randomly decide whether to reconnect one end of the edge to a randomly selected node in the network. The (independent) probability of this happening to any edge is p_r . When all edges have been considered, the procedure stops and the resulting network is our starting point. When $p_r = 0$, the structure is the initial (minimally connected) caveman graph. When $p_r = 0.1$, the structure can be described as a semi-isolated caveman graph. When $p_r = 1$, the rewired network that supports social learning is a uniform random graph in which edges exist with identical and independent probability. All three networks have the same density, yet very different architectures.

We saw in the first experiment that a random uniform network of intermediate density allowed maximal learning accuracy over a large set of network sizes (see Figure 1). The second experiment further refines these results by exploring the impact of network architecture, particularly focusing on

local clustering and cliquishness. Each clique will rapidly gather information and converge to one alternative, but as was noted in the first experiment, dense graphs can fail to identify the superior alternative. Decreasing intra-clique density can partly remedy the problem, delaying homogenization until enough time has permitted local exploration to identify the superior arm. However, in the type of structures we consider, decreasing intra-clique density comes at the expense of increased inter-clique connectivity, which can result in speeding up homogenization. Depending on which effect dominates, different conclusions might obtain.




Figure 5: The accuracy of learning after 2,000 periods in communities of various cave sizes and rewiring probabilities p_r ; each curve is for one cave size $n_c = 3, 6, 8, 12, 16, 24, 48$

Figure 5 depicts learning performance after 2,000 periods as it responds to the rewiring probability p_r for different cave sizes, replicating Fang et al. (2010) with social Bayesian learning on two-arm bandits. The x-axis is logarithmic, to increase legibility.

We first observe that the results of Fang et al. carry over to the context of Bayesian learning: for any cave size, there is a strong effect of the first very few rewired links on performance, but the effect of additional links peters out rapidly past a rewiring probability of $p_r = 0.1$. More interestingly, for $n_c = 12, 16$ and 24 , an interior maximum obtains for a rewiring probability of 0.1. At this value of the rewiring probability, network structure corresponds to what Fang et al. characterize as “semi-isolated” caves. Coming back to the finding of the first experiment that intermediate network density allowed maximal learning accuracy, we observe now that intermediate cave sizes (of 12, 16 and 24 agents) outperform caves of both smaller and larger sizes (with the exception of the largest

cave size of 48 when rewiring is extreme) and that for such caves, maximal performance is obtained when the rewiring probability is small but non zero. We therefore conclude that although density plays a major role in determining the learning accuracy of a community, for a given density there are subtle architectural differences that also impact the quality and speed of learning.

Figure 6 represents learning performance as a function of cave size n_c for different levels of the rewiring probability p_r (the performance metric is the same as in the previous figure, but cave size and rewiring probability are swapped). As cave size perfectly correlates with the density of the entire network, the pattern echoes the one depicted in Figure 1: increasing cave size (and thus density in the overall network) increases for any rewiring intensity, up to the point where the less-is-more effect kicks in. Past this point, aggregate performance degrades as cave size increases further.




Figure 6: The accuracy of learning after 2,000 periods in communities of various cave sizes and rewiring probabilities p_r ; each curve is for one rewiring probability.

Discussion

Our simulations suggest the following refinements to the existing findings on the less-is-more effect. First, in a setup similar to that used by Zollman (2010), we did not observe a clear less-is-more effect in communities of larger size than those considered by Zollman. The effect is present over part of the density spectrum, but intermediate levels of density are necessary to strike a satisfactory balance between short term exploration and long term diffusion.

When studying semi-isolated groups of agents, we were able to find a set of conditions under which collective per-

formance worsens as density increases, as suggested by the less-is-more effect (Figure 6). The effect of cave size (which is equivalent to density in the family of networks we consider) is non-monotonic, however, and at lower network densities, more seems to be more, i.e. increases in density improve collective performance, up to a limit.

Regarding the work of Fang et al. (2010), we find that to some extent, their results carry over to a context of Bayesian learning⁷: at first, rewired links dramatically improve performance, but as additional edges are rewired, collective performance plateaus and then declines.

All in all, our study suggests several refinements to the less-is-more hypothesis. They should be seen as call for modesty: In our view, our findings, combined with mixed evidence from empirical studies (Mason et al., 2008; Mason & Watts, 2012; Shore et al., 2015) suggests that the story behind the less-is-more effect is more complicated than previously assumed, and all of the mechanisms driving the effect are not yet fully understood. How connectivity and transient diversity influence distributed learning in a population of agents depends on the size of the community as well as its detailed patterns of connectivity. Our findings indicate, however, that semi-isolated clusters help to strike a balance fostering epistemically beneficial transient diversity.

Acknowledgments

We thank the four COGSCI reviewers for useful feedback. Samuli Reijula's research is supported by the Research Council of Finland.

References

- Bala, V., & Goyal, S. (1998). Learning from neighbours. *The review of economic studies*, 65(3), 595–621.
- Barkoczi, D., & Galesic, M. (2016). Social learning strategies modify the effect of network structure on group performance. *Nature communications*, 7(1), 13109.
- Berry, D. A., & Fristedt, B. (1985). *Bandit problems*. Dordrecht: Springer Netherlands.
- Fang, C., Lee, J., & Schilling, M. A. (2010, June). Balancing Exploration and Exploitation Through Structural Design: The Isolation of Subgroups and Organizational Learning. *Organization Science*, 21(3), 625–642.
- Grimm, V., Railsback, S. F., Vincenot, C. E., Berger, U., Gallagher, C., DeAngelis, D. L., ... Groeneveld, J. (2020). The ODD protocol for describing agent-based and other simulation models: A second update to improve clarity, replication, and structural realism. *Journal of Artificial Societies and Social Simulation*, 23(2).
- Kummerfeld, E., & Zollman, K. J. S. (2016, December). Conservatism and the Scientific State of Nature. *The British Journal for the Philosophy of Science*, 67(4), 1057–1076.

⁷Fang and coauthors use a different learning mechanism, one based on the March (1991) model of socialization.

- Lazer, D., & Friedman, A. (2007). The network structure of exploration and exploitation. *Administrative Science Quarterly*, 52(4), 667–694.
- March, J. G. (1991). Exploration and exploitation in organizational learning. *Organization science*, 2(1), 71–87.
- Mason, W., Jones, A., & Goldstone, R. L. (2008). Propagation of innovations in networked groups. *Journal of Experimental Psychology: General*, 137(3), 422.
- Mason, W., & Watts, D. J. (2012). Collaborative learning in networks. *Proceedings of the National Academy of Sciences*, 109(3), 764–769.
- Mehlhorn, K., Newell, B. R., Todd, P. M., Lee, M. D., Morgan, K., Braithwaite, V. A., ... Gonzalez, C. (2015). Unpacking the exploration–exploitation tradeoff: A synthesis of human and animal literatures. *Decision*, 2(3), 191.
- Rolin, K., Koskinen, I., Kuorikoski, J., & Reijula, S. (2023, July). Social and cognitive diversity in science: introduction. *Synthese*, 202(2).
- Shore, J., Bernstein, E., & Lazer, D. (2015, October). Facts and Figuring: An Experimental Investigation of Network Structure and Performance in Information and Solution Spaces. *Organization Science*, 26(5), 1432–1446.
- Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction*. MIT press.
- Watts, D. J. (1999). Networks, Dynamics, and the Small-World Phenomenon. *American Journal of Sociology*, 105(2), 493–527.
- Yahosseini, K. S., Reijula, S., Molleman, L., & Moussaid, M. (2018). Social information can undermine individual performance in exploration-exploitation tasks. In *COGSCI 2018: 40th Annual Conference of the Cognitive Science Society* (pp. 2470–2475).
- Zollman, K. J. (2010). The Epistemic Benefit of Transient Diversity. *Erkenntnis*, 72(1), 17–35.