DOI: 10.1017/psa.2026.10275
This is a manuscript accepted for publication in *Philosophy of Science*.
This version may be subject to change during the production process.

Theory Choice in Epistemic Networks: Five ways to avoid premature convergence

Nicolas Jonard[1], Samuli Reijula[2*] and Luigi Marengo[3]

[1]Department of Economics and Management, University of Luxembourg, Luxembourg, L1359, Luxembourg
[2]Theoretical Philosophy, PO Box 24 , 00014 University of Helsinki, Finland
[3]Department of Business and Management, LUISS University, Rome, 00197, Italy
[*]Corresponding author: Samuli Reijula; samuli.reijula@helsinki.fi

Abstract

In this article, we study difficult theory-choice situations, where division of cognitive labor is needed. Network epistemology models suggest that reducing connectivity is needed to prevent premature convergence on bad theories. We compare how network density, community size, strength of prior beliefs, adaptive learning methods, and weak ties influence epistemic outcomes, and show that reducing connectivity is only one possible way to improve collective epistemic accuracy. Our findings suggest that gains in accuracy often come at a high cost in resources used, which should be considered when results from network epistemology models are used in applied settings.

1. Introduction

Once there was a Very Difficult Problem. Perhaps the problem concerned discovering the structure of a very important molecule, but not knowing which of two promising research paths to follow. Or maybe it was about choosing between two medical treatments, where in experimental trials, the difference between the correct and incorrect treatment was difficult to detect.

How should we proceed in such difficult theory-choice situations? This article suggests that if the problem is hard enough to require division of cognitive labor between a community of scientists, several aspects of cognitive coordination will influence the accuracy and efficiency of the epistemic community: Making the epistemic community large enough, if feasible, will lead to quick and accurate choices. So will introduction of adaptive decision rules for scientist agents: even if scientists follow relatively simple heuristics for balancing exploration and exploitation in their experimental activity, the true theory will be discovered efficiently.

Manipulating connectivity and strength of prior beliefs, as suggested by Zollman (2010), are two other ways of ensuring sufficient exploration of alternative theoretical options.[1] We show, however, that the less-is-more effect reported by Zollman should

[1]For overviews of the transient diversity phenomenon, see Wu and O'Connor 2023; Smaldino et al. 2024.

https://doi.org/10.1017/psa.2026.10275 Published online by Cambridge University Press

be primarily understood as an asymptotic finding: Our simulations confirm that as long as a community of scientists remains connected (even through weak links), reducing connectivity leads to improvements in accuracy. However, if resource constraints, such as cost of experimentation or time, are introduced, extremely low levels of connectivity are an inefficient way to improve accuracy. For all practical purposes, there comes a point where very little connectivity is too little.

1.1. A model of collective theory choice

As several authors have suggested, a simple way to model theory choice in the context of the Very Difficult Problem is to assume that the scientist is confronted with the task of finding the better arm of a two-armed Bernoulli bandit. Following Zollman (2010) and others[2] we assume that each arm $i \in \{A, B\}$ produces a reward of 1 with success probability p_i and 0 with complementary probability $1 - p_i$. Specifically, we assume $p_A = 0.5$ and $p_B = p_A - 0.01 = 0.49$, i.e., the difference in the probability of success is small. The success probabilities (p_A, p_B) are unknown to the scientist but fixed over time.

The scientist can collect evidence on the success probability of each option by running experiments, which amounts to sampling the arms of the bandit. As is common in the literature, we assume the scientist's prior beliefs about p_i, the success probability of arm i, to be Beta-distributed over $[0, 1]$ with parameters $\alpha_i, \beta_i > 0$. In particular, for each option i, the prior probability density function f of p_i is

$$f(p_i) = \frac{\Gamma(\alpha_i + \beta_i)}{\Gamma(\alpha_i)\gamma(\beta_i)} p_i^{\alpha_i - 1}(1 - p_i)^{\beta_i - 1},$$

where Γ denotes the Gamma function. The mean success probability of arm i thus obtains as the mathematical expectation of the Beta-distributed prior beliefs, as

$$\frac{\alpha_i}{\alpha_i + \beta_i}.$$

When the agents follow what is called the greedy or *myopic decision rule*, the mean success probability (given the belief distribution held by the scientist) is assumed to govern arm selection. When the scientist chooses between A and B, they compare the mean success probabilities of the two arms and choose the arm with the larger mean success probability. A tie is broken by tossing a fair coin.

Once an arm is selected, it generates a number of observations, and the distribution of beliefs gets updated according to Bayes' rule. This is where Beta distributions are particularly convenient: when the prior is Beta-distributed, the posterior is also Beta-distributed with parameters easily updated. Specifically, a number n_i of trials (arm pulls) on arm i produces a binomially distributed random number of successes, $X_i \sim \mathrm{B}(n_i, p_i)$, with $0 \leq X_i \leq n_i$. The posterior mean success probability of arm i is obtained as

$$\frac{\alpha_i + X_i}{\alpha_i + \beta_i + n_i}. \tag{1}$$

[2]See for instance Bala and Goyal (1998); Rosenstock et al. (2017); Frey and Šešelja (2020); Fazelpour and Steel (2022); Wu (2023); Michelini et al. (2023).

https://doi.org/10.1017/psa.2026.10275 Published online by Cambridge University Press

1.2. A benchmark

Let's first try to get a sense of how many experiments would ideally be needed to solve the Very Difficult Problem.

Rather than following a specific reinforcement learning rule to select an arm to pull in each round, we ask how many pulls the scientist should perform on the two arms to reduce uncertainty in a way that makes them capable of accurately distinguishing the superior arm from the inferior one. We assume that an acceptable level of epistemic performance in this choice problem is an accuracy of 90%. This will require a number of pulls on each arm such that the posterior of the mathematical expectation of the probability of success of the arm A exceeds that of arm B with probability at least 0.9. A moment's reflection suggests that it is the size of the sample drawn from the less sampled arm which will constrain the accuracy with which the best arm is identified.

Letting n_A (resp. n_B) be the number of pulls on arm A (resp. B), the probability that the posterior mean success probability of arm A is larger than that of arm B, which we informally denote $p_{A>B}$, is written

$$\Pr\left\{\frac{\alpha_A + X_A}{\alpha_A + \beta_A + n_A} > \frac{\alpha_B + X_B}{\alpha_B + \beta_B + n_B}\right\} + \frac{1}{2}\Pr\left\{\frac{\alpha_A + X_A}{\alpha_A + \beta_A + n_A} = \frac{\alpha_B + X_B}{\alpha_B + \beta_B + n_B}\right\}.$$

Defining

$$\theta(x) = \frac{\alpha_A + \beta_A + n_A}{\alpha_B + \beta_B + n_B} \times (\alpha_B + x) - \alpha_A,$$

and conditioning on the possible values of the (binomially distributed) random variable X_B, the probability of identifying the superior arm, $p_{A>B}$, is finally obtained as

$$\begin{aligned}
&\Pr\{X_A > \theta(X_B)\} + \frac{1}{2}\Pr\{X_A = \theta(X_B)\} \\
&= \sum_{x_B=0}^{x_B=n_B} (1 - \Pr\{X_A \leq \theta(x_B)\}) \Pr\{X_B = x_B\} \\
&+ \frac{1}{2} \sum_{x_B=0}^{x_B=n_B} \Pr\{X_A = \theta(x_B)\} \Pr\{X_B = x_B\}.
\end{aligned}$$

As we know the density and distribution functions of binomial distributions, this probability is easily calculated. In Figure 1 we represent the probability of identifying the superior arm, $p_{A>B}$, as a function of n_A and n_B, the number of pulls performed on arms A and B. The parameters of the initial distributions of beliefs are set to $\alpha_A = \beta_A = \alpha_B = \beta_B = 1$, so that the initial priors are uniformly distributed over $[0, 1]$. The contour lines help identify our epistemic performance threshold of 90%.

Reading the contour lines, one sees that a 0.9 probability of identifying the superior arm requires roughly $10^{3.9} \approx 8,000$ trials with each arm (increasing to a 0.95 probability would require as many as $10^{4.1} \approx 12,500$ trials per arm). This can be considered as our informal benchmark, the threshold amount of information needed to reach the acceptable level of accuracy. Observe that there is little substitutability across the number of pulls performed on each arm: raising the probability of superior arm identification

https://doi.org/10.1017/psa.2026.10275 Published online by Cambridge University Press

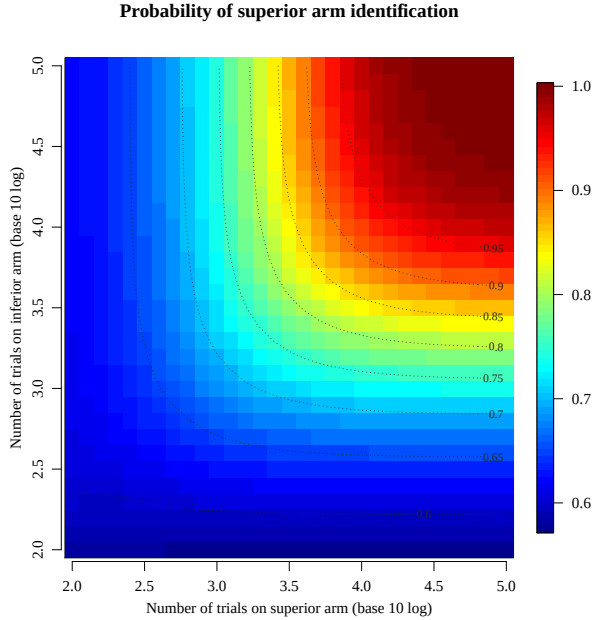


Figure 1. The probability of identifying the superior arm, $p_{A>B}$, as a function of the number of trials, $n_i, i \in \{A, B\}$, performed with each arm, with this number between 100 and 100,000. The axis are base 10 logarithmic, the probability is depicted as a heatmap with contour lines, the initial priors are uniformly distributed, with $\alpha_A = \beta_A = \alpha_B = \beta_B = 1$.

demands that both arms are tried more often. Focusing only on one arm (moving vertically or horizontally in Figure 1) can only increase the probability of identifying the superior arm up to a point. If, say, 10^3 trials have been performed with the superior arm, performing an increasing number of trials on the inferior arm (moving upward on the vertical line that runs through 3 on the horizontal axis) is not going to improve accuracy much beyond 0.7, and certainly not up to 0.9. This confirms that it is the size of the smaller sample that constrains the accuracy with which the best alternative is identified.

With this benchmark at hand, we now turn to the different ways of organizing an epistemic community that allow it to draw the required number of experiments from the less sampled option. In Section 2, we show how the problem can be solved by the sheer size of the epistemic community, but why a densely connected network of insufficient size fails. Furthermore, dense connections between network members lead to epistemic harm only under the myopic rule, otherwise more connectivity does no harm. Section 3 shows that strong prior beliefs solve the problem of premature convergence, but at a considerable price to efficiency. Section 4 pushes the less-is-more effect discovered by Zollman to its limit: We study epistemic networks connected by weak links, where information between scientists is transmitted at specific time intervals only. It seems that less connectivity is better, but again, only when resource constraints are not taken into account.

https://doi.org/10.1017/psa.2026.10275 Published online by Cambridge University Press

2. Crowdsourcing the problem: large networks

8,000 experiments is a lot. Assuming a non-zero duration or cost for running one individual experiment results in a sequence that no scientists could run on their own. That's how it should be. After all, we're assuming the problem is very difficult. In the philosophy of science, there is a broad agreement that science is a deeply collective enterprise that relies on division of epistemic labor between the members of a scientific community. How should that community be organized?

Existing literature on epistemic networks suggests that simply giving all agents in a community access to each other's experimental evidence will lead to worse outcomes than in situations where information flows are more restricted (Zollman, 2010). That is surprising: given the framing of the problem just presented, it seems that more connectivity increases the amount of information available to the members of the community, and in the absence of bias or egotistic motives of scientists (as a contrast, see sullied agents by Kitcher 1990), more information should be better.

To understand why limiting information flow may sometimes be necessary, we must look in more detail into how cognitive coordination plays out in this case. We borrow from Zollman's work our assumptions about how to set up the socially distributed version of the theory choice task (for a systematic presentation of the details of the model, see the Supplementary Information document).

Paralleling the case of a single scientist, we assume that at each time step each scientist j in a community $N = \{1, \ldots, j, \ldots, m\}$ must decide which theory to work on next. When agents follow the myopic decision rule, they always sample the arm that they think is better. Despite being very simple, the assumption does plausibly capture an aspect of the dynamics of collective problem solving. If we assume that no agent knows when the experimentation will end and that their current action stands for their overall decision of which theoretical approach to root for, it can be seen as rational for an agent to always choose the option that currently appears better in light of their beliefs.[3]

Figure 2 summarizes the simulation results from communities of different size and density. There are 6 values for density in $\{0, 0.1, 0.25, 0.5, 0.75, 1\}$ and 6 values for network size in $\{3, 5, 10, 25, 50, 100\}$. For each density value crossed with each network size value, we perform 500 independent runs of 5,000 rounds, with $\tau = 100$ pulls for each agent per round. In each parameter configuration, we set $\alpha_{j,A} = \beta_{j,A} = \alpha_{j,B} = \beta_{j,B} = 1$ for each agent j so that all initial distributions of beliefs are uniformly distributed over $[0, 1]$. This assumption implies that all agents toss a coin in the first round, and approximately half of them will choose each arm.[4]

The left panel of Figure 2 essentially replicates and extends findings from Zollman (2010) and Rosenstock et al. (2017). The negative relationship between the proportion of successful runs and network density that is observable for all community sizes is what has been called the *Zollman effect*: less connectivity leads to better collective accuracy.

A typical explanation for the poor performance of dense networks is to note that they suffer from premature convergence on suboptimal options, a phenomenon observed

[3]That said, the notion of 'rooting for an an alternative' is a little unclear here. Why would it not be enough of a goal for scientists simply to get their beliefs right, $E[f_A] > E[f_B]$? In other words, if the purpose of inquiry is cognitive, i.e. it only aims at true belief, the sampling action that an agent performs at any given time step should be irrelevant to that cognitive goal.

[4]Zollman (2010) makes a slightly different assumption, assuming αs and βs are sampled uniformly from $[0, 4]$. The discussion in Section 3 shows that both assumptions result in very similar outcomes.

https://doi.org/10.1017/psa.2026.10275 Published online by Cambridge University Press

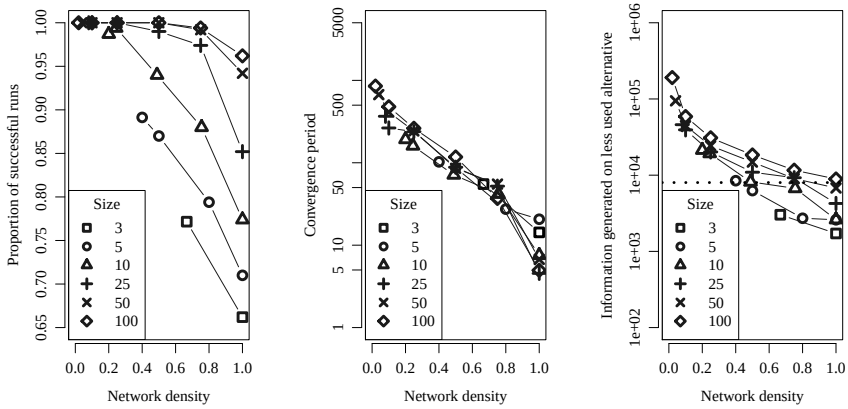


Figure 2. The proportion of successful runs after 5,000 rounds (left panel), the average value of the round in which one arm was definitively abandoned (middle panel) and the total number of pulls on the less sampled arm (right panel) in networks of various sizes and densities, when each round consists of $\tau = 100$ trials. Each curve is for one network size, averaging 500 independent runs.

across a broad range of models in different contexts (e.g., Lazer and Friedman (2007); March (1991); Fang et al. (2010)). But how does it work in this model setup, precisely? The explanation turns out be significant, because by better understanding the conditions of premature convergence, we see that high connectivity does not always lead to bad epistemic outcomes.

We next look into the micro-level events that can give rise to premature convergence in dense networks. First, however, an observation on the framing of the Zollman effect: The Very Difficult Problem is indeed very difficult, as we noted in the beginning. Therefore, the fact that a dense network of size 10 achieves an accuracy of only around 0.8 should not be seen as disastrous. As our benchmark simulation suggests, reaching such an accuracy requires running at least roughly 3000 experiments. With a sample size in the hundreds, a network's performance amounts to basically a 50-50 guess between two almost equally good options. In contrast, what is really surprising about the Zollman effect is how amazingly accurate the sparse networks are. They reach almost full accuracy in this difficult task. How can *that* be?

2.1. The failure of fully connected networks

We assume that for any network size m, each agent $j \in \{1, \ldots, m\}$ has identical uniform beta priors captured by $\alpha_{j,A} = \beta_{j,A} = \alpha_{j,B} = \beta_{j,B} = 1$. Given that each agent tosses a fair coin to pick an arm and pulls it τ times to collect a sample of observations, in the first round the population will collectively draw roughly equal-sized samples from both bandit arms. In the fully connected network, all samples are then pooled, and the resulting (large) sample is made available to all. After updating, all agents now have exactly the same posterior distribution, and unless the community has collectively generated a draw, all scientists will pull the same arm the next round, and abandon the other one. This will result in massive evidence being generated again, but only for one arm. Most likely, agents will continue to choose that arm indefinitely. (It is possible that a

https://doi.org/10.1017/psa.2026.10275 Published online by Cambridge University Press

downward-biased sequence of pulls causes all agents to switch to the other arm, but this becomes increasingly unlikely as evidence accumulates.)

In other words, if the initial sample from arm A suggests an estimate of p_A significantly less than 0.49, and the sample from arm B is not similarly negatively biased, it may be that arm A is never sampled again: Every agent in the population believes that B is the better option, and as they continue drawing samples from option B, their estimate of p_B becomes ever more accurate. However, as we showed in Section 1, this is not sufficient to ensure epistemic accuracy. Sampling the better arm has a decreasing marginal contribution to epistemic precision: the precision of judgment is limited by the inaccuracy of the belief distribution about the less sampled arm A.[5] Therefore, after the first period, a fully connected epistemic community essentially acts as a single myopic scientist, quickly abandoning exploration.

That said, the higher the number of agents, the larger the amount of information collected, and the higher the epistemic accuracy we should expect, based on Figure 1. Although this may sound at odds with the Zollman effect, focusing on the rightmost part of the horizontal axis in the left panel of Figure 1 clarifies the apparent paradox. The performance of the complete network (density of 1) increases as its size increases, moving from 0.65 for a community of 3 scientists to more than 0.95 for a community of size 100. Why is that?

Looking at the middle panel in Figure 2 reveals that larger complete networks stop exploring on average earlier than smaller ones, and most networks (of size 10 to 100) take between 5 and 10 rounds to converge. However, more important than the convergence period is the amount of information generated on the less sampled arm, which is depicted on the right panel of Figure 1. Unlike the dense but small networks, networks of size 50 and 100 generate a sufficient amount of information already during the first rounds, and land above the threshold amount of information needed to reach the acceptable level of epistemic precision identified earlier (indicated with the dashed line).

What this suggests is that *one way to solve the Very Difficult Problem is to crowdsource it, to hand it over to a large and densely connected community of scientists*. If such a large population of scientists is available, such crowdsourcing can deliver desired accuracy very quickly, only after a couple of rounds of experimentation.

2.2. Less is more only for myopic agents

The less-is-more phenomenon characteristic of the Zollman effect should therefore not be understood not in terms of information quantity (where too little can definitely be too little, as the small-sized, fully connected networks indicate), but exclusively in terms of network density. *Capacity for information generation* is the crucial aspect distinguishing dense and sparse networks: given network size, sparse networks generate more information (on the less sampled arm) than denser networks, and thus the former outperform the latter. This is seen in Figure 2, where on any curve, moving left in terms of density means moving up in terms of accuracy (left panel), through the effect of moving up in terms of information generation (right panel).

[5]See the generalized hot stove effect in Denrell (2024) for an explanation of essentially the same phenomenon in a general learning context.

https://doi.org/10.1017/psa.2026.10275 Published online by Cambridge University Press

But what creates the connection between network density and information generation? To answer this question, we have to start from the dual nature of the *sequential sampling task* (Denrell and March, 2001) faced by the simulated scientists. In such tasks, agents simultaneously pursue the dual objectives of collecting payoff (for which pulling the superior arm is essential) and collecting information (for which pulling also the other arm is necessary).

In the experiments we have discussed so far, the tension between the two conflicting objectives is resolved by the myopic decision rule, which always prioritizes payoff generation. As Figure 3 shows, if we relax the rule, and allow our agents to balance exploitation and exploration more adaptively, e.g. as in *Thompson sampling*, the less-is-more effect of network density disappears.[6]

Whereas the myopic rule only tracks the expectations of the current belief distributions, Thompson sampling (also called probability matching) considers random deviates generated from the current belief distributions (Thompson, 1933). A deviate is generated for each arm using the current Beta distribution of beliefs of that arm, and the arm with higher-valued deviate is selected. A sample of size τ is then generated for the selected arm and the belief distribution updated as previously described.

This approach significantly improves over myopic (or its variant, ε-greedy) selection, as it takes into account the entire distribution (and in particular, variance) of each arm's belief distribution. In Thompson sampling, both expected payoff and uncertainty regarding an option influence the sampling decision.

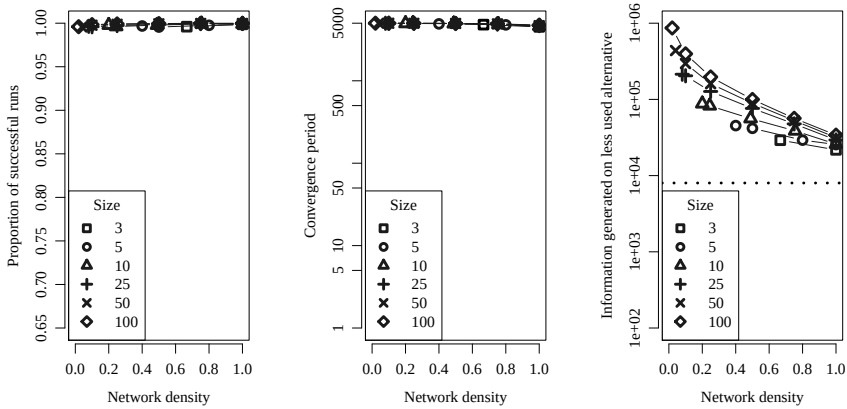


Figure 3. A population of agents using Thompson sampling. The proportion of successful runs after 5,000 rounds (left panel), the average value of the round in which one arm was definitively abandoned (middle panel) and the total number of pulls on the less sampled arm (right panel) in networks of various sizes and densities, when each round consists of 100 trials; each curve is for one network size, averaging 500 independent runs

Figure 3 parallels Figure 2, summarizing the simulation results from a population of agents relying on Thompson sampling. Here, both the value and the cost of Thompson

[6]See Kummerfeld and Zollman (2016) for analysis of the model with fixed exploration rates, i.e. ε-greedy sampling.

https://doi.org/10.1017/psa.2026.10275 Published online by Cambridge University Press

sampling are manifest. After 5,000 periods, regardless of their size and density, all networks have detected the superior arm (left panel) but convergence of action has not occurred (middle panel). The amount of information generated, though decreasing with density, vastly exceeds the threshold amount of information we identified in the first section: by a factor at least 5 for small, dense networks, up to factor 100 for large, sparse networks.

3. Opinionated scientists

Zollman suggests that one way to avoid premature convergence is to give agents extreme initial priors, which will make them explore for longer. The structure of the model suggests that strong prior beliefs are particularly important in large networks. We have assumed that the sample collected by each agent in one time step is $\tau = 100$. Then at every round, 100 pieces of evidence from the agent's own experiments, plus the evidence coming from direct network neighbors get added the agent's total experimental evidence (See SI, Section 2, for a visual illustration). Equation 1 shows that in order for the agent's priors to continue to have any effect on the expectation, α_A and β_A (resp. for B) must be commensurate with the amount of new evidence reaching the agent in one time step.

Figure 4 again parallels Figure 2, now for opinionated agents. That is, the agents' initial priors are Beta-distributed with parameters sampled between 1 and $\tau \times (1 + \overline{d})$, where $\overline{d}$ is the average number of network neighbors. Observe that average degree and network density are directly proportional, since density is equal to

$$\frac{\sum_{j=1}^{j=N} d_j}{N(N-1)} = \frac{1}{N-1}\frac{\sum_{j=1}^{j=N} d_j}{N} = \frac{\overline{d}}{N-1},$$

so that denser networks are associated with more extreme beliefs by construction.

The benefit of strong initial priors is very clear and the less-is-more effect of network density vanishes (the curves on the left panel show essentially no pattern, except for the smallest network). Larger networks do better than smaller ones, to an extent that is roughly independent from density. Exploration has stopped (middle panel) much later than in the original experiment reported in Figure 2, with convergence postponed to at least 200 periods for small networks, and up to 4,000 periods in large networks. Right panel shows why the less-is-more effect has vanished: For any network size, the amount of information generated is roughly independent of network density: the more extreme priors attached with denser networks have forced more exploration, yielding roughly constant information generation (flat curves in the right panel). The amount of information generated, however, vastly exceeds the threshold amount of information we identified in the first section: by a factor at least 2 for small, dense networks, up to factor 100 for large, sparse networks.

If we assume that information generation is costly (i.e. non-zero cost to an experiment), we may conclude that *strong priors are an effective but inefficient way to identify the superior arm.*

https://doi.org/10.1017/psa.2026.10275 Published online by Cambridge University Press

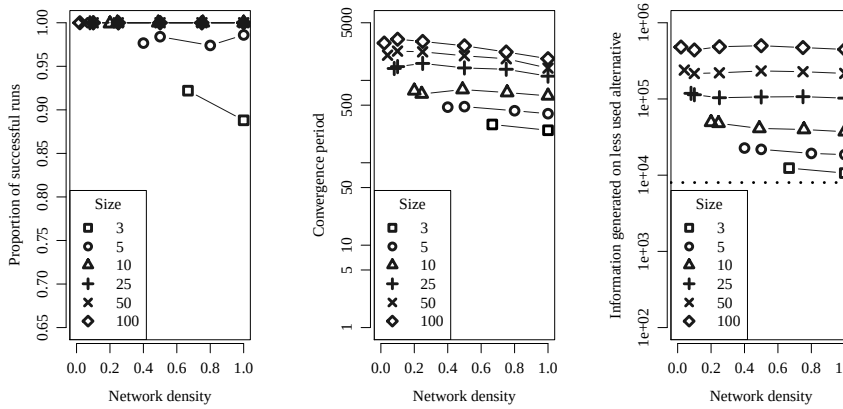


Figure 4. The proportion of successful runs after 5,000 rounds (left panel), the average value of the round in which one arm was definitively abandoned (middle panel) and the total number of pulls on the less sampled arm (right panel) in networks of various sizes and densities, when each round consists of 100 trials; each curve is for one network size, averaging 500 independent runs; the initial priors are strong, in the sense of being Beta-distributed with parameters sampled between 1 and $\tau \times (1 + \overline{d})$, where $\overline{d}$ is the average number of connections per agent.

4. Weak ties

Sparsely connected networks perform remarkably well in solving the Very Difficult Problem. But how little connectivity is too little? Surely, in order for the less-is-more effect to be of theoretical interest, there must be a minimum level of connectivity below which collective epistemic performance deteriorates.

Note that even in the cycle network where every agent has only two direct neighbors, the amount of social information utilized by an agent exceeds the amount information they generate themselves (by a ratio of two to one). To limit the transmission of social information, strong network links can be replaced with weak ties that activate only intermittently. Findings by Smaldino et al. (2024) and Bernstein et al. (2018) suggest that epistemic outcomes can be improved if the communication rate between neighbors is reduced.

In order to study situations where an agent's information mainly comes from their own experiments, assume that each agent samples on their own for $s - 1$ rounds, and every s rounds the epistemic network becomes activated, allowing neighbors to share the evidence from the current round. After which the cycle between $s - 1$ rounds of solitary experimentation and 1 round of social learning repeats. Social learning is intermittent, with intermittence period $s \geq 1$. In this scenario, the ratio of individual to social information is s over the agent's number of neighbors.

Starting with the left panel in Figure 5, first observe that the curve corresponding to an intermittence period of 1 is exactly the same as the one on the left panel in Figure 2, when network size is 25. The less-is-more effect is clearly at work, with accuracy falling with density.

Intermittent social learning dramatically changes the outcome. Although increasing the intermittence period increases information generation (right panel) and postpones

https://doi.org/10.1017/psa.2026.10275 Published online by Cambridge University Press

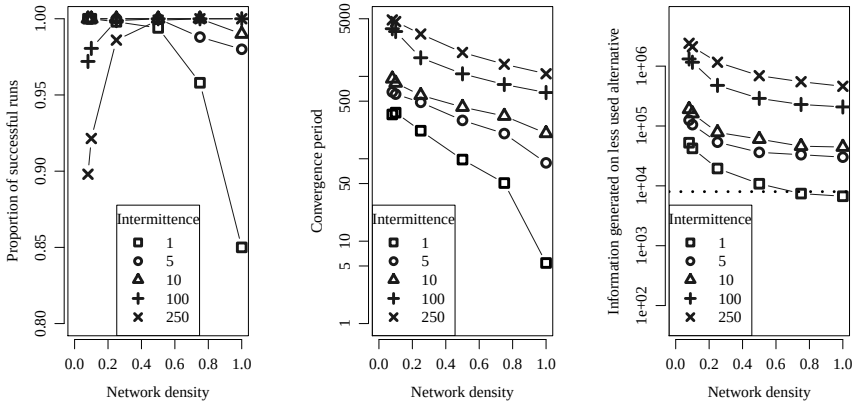


Figure 5. The proportion of successful runs after 5,000 rounds (left panel), the average value of the round in which one arm was definitively abandoned (middle panel) and the total number of pulls on the less sampled arm (right panel) in a network of 25 agents and varying density, when each round consists of $\tau = 100$ trials, but social learning is intermittent, with period s varying between 1 (the situation considered so far) and 250; each curve is for one level of s, averaging 500 independent runs.

convergence (middle panel) for any value of density, its effect on accuracy is ambiguous. Increasing the intermittence period s benefits fully connected networks but is very detrimental to minimally connected networks. Intermittence addresses the shortcomings of fully connected networks by delaying convergence that results from premature termination of exploration, but results in over-exploration in sparse networks in which convergence is already slow. Consequently, minimally connected networks have not converged after 5,000 periods and the generation of several millions of observations on each arm. Here, the trade-off between speed and accuracy is manifest. In this respect, intermediate intermittence values ($s = 10$ in the figure) seem to achieve the best outcome regardless of the density of the network.

If information generation is costly, *intermittent social learning is an effective and efficient way to identify the superior arm in dense networks, less so in sparser ones*.

5. Conclusions

We have argued that when it comes to acquiring information in a collective theory choice situation, more is more. In particular, the amount of information collected on the less sampled option is the factor limiting accuracy in epistemic networks. Understanding the mechanisms through which this factor drives aggregate outcomes makes it clear that reducing network connectivity is only one way to increase accuracy. The sheer size of the network, adaptive learning rules, strong prior beliefs, and intermittently occurring social learning are among other ways of ensuring that all theoretical options are given a fair hearing.

We concluded each section with a comment on the efficiency of the aforementioned ways of reaching high epistemic accuracy. Not surprisingly, increasing the rationality of learning by replacing the myopic rule with a more adaptive option appears most efficient. Strong prior beliefs, in contrast, lead to large amounts of experiments

https://doi.org/10.1017/psa.2026.10275 Published online by Cambridge University Press

being conducted, and should be seen a costly way of preventing premature convergence. Likewise, as we showed in Section 4, if learning from network neighbors is too infrequent, convergence on the correct option is delayed.

We believe that in order to translate findings from network epistemology models to more general contexts, both in the philosophy of science and beyond, considerations of efficiency are necessary: For many purposes, it is not enough to know that an epistemic community will eventually converge on the true option, as time approaches infinity: Note that in the model setting introduced by Zollman (2010), and also studied in this article, the following holds: As long as at all times there is a non-zero probability that the less-studied arm gets sampled by even a single agent connected to the rest of the network (by at least one link), the rest of the population will eventually converge on the better arm. This condition can be satisfied in various ways, for example, by introducing a single ε-greedy agent in an otherwise myopic population. Consensus on the true option would be reached, but the process would be highly inefficient.

Hence the need to phrase questions of collective theory choice in light of resource constraints, a project we hope to have initiated in this article.

Acknowledgments We thank the participants at PSA 2024 and at TINT Brown Bag seminar (University of Helsinki) for helpful comments and suggestions.

Declarations of competing interests None to declare

Funding Information Samuli Reijula's research was funded by the Research Council of Finland (#332686).

References

Bala, V. and Goyal, S. (1998) Learning from neighbours. *The review of economic studies*. *65*(3), 595–621.

Bernstein, E., Shore, J. and Lazer, D. (2018) How intermittent breaks in interaction improve collective intelligence. *Proceedings of the National Academy of Sciences*. *115*(35), 8734–8739. 10.1073/pnas.1802407115.

Denrell, J. (2024) Adaptive Sampling Policies Imply Biased Beliefs: A Generalization of the Hot Stove Effect. 10.48550/arXiv.2404.02591. arXiv:2404.02591 [cs].

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

Fang, C., Lee, J. and Schilling, M. A. (2010) Balancing Exploration and Exploitation Through Structural Design: The Isolation of Subgroups and Organizational Learning. *Organization Science*. *21*(3), 625–642. 10.1287/orsc.1090.0468.

Fazelpour, S. and Steel, D. (2022) Diversity, trust, and conformity: A simulation study. *Philosophy of Science*. *89*(2), 209–231.

Frey, D. and Šešelja, D. (2020) Robustness and Idealizations in Agent-Based Models of Scientific Interaction. *The British Journal for the Philosophy of Science*. *71*(4), 1411–1437. 10.1093/bjps/axy039.

Kitcher, P. (1990) The Division of Cognitive Labor. *Journal of Philosophy*. *87*(1), 5–22. 00241.

Kummerfeld, E. and Zollman, K. J. S. (2016) Conservatism and the Scientific State of Nature. *The British Journal for the Philosophy of Science*. *67*(4), 1057–1076.

Lazer, D. and 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.

Michelini, M., Osorio, J., Houkes, W., Šešelja, D. and Straßer, C. (2023) Scientific Disagreements and the Diagnosticity of Evidence: How Too Much Data May Lead to Polarization. *Journal of Artificial Societies and Social Simulation*. *26*(4), 5.

https://doi.org/10.1017/psa.2026.10275 Published online by Cambridge University Press

Rosenstock, S., Bruner, J. and O'Connor, C. (2017) In Epistemic Networks, Is Less Really More? *Philosophy of Science*. *84*(2), 234–252.

Smaldino, P. E., Moser, C., Pérez Velilla, A. and Werling, M. (2024) Maintaining Transient Diversity Is a General Principle for Improving Collective Problem Solving. *Perspectives on Psychological Science*. *19*(2), 454–464. 10.1177/17456916231180100.

Thompson, W. R. (1933) On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*. *25*(3/4), 285–294.

Wu, J. (2023) Epistemic advantage on the margin: A network standpoint Epistemology. *Philosophy and Phenomenological Research*. *106*(3), 755–777. 10.1111/phpr.12895.

Wu, J. and O'Connor, C. (2023) How should we promote transient diversity in science? *Synthese*. *201*(2), 37. 10.1007/s11229-023-04037-1.

Zollman, K. J. (2010) The Epistemic Benefit of Transient Diversity. *Erkenntnis*. *72*(1), 17–35.

https://doi.org/10.1017/psa.2026.10275 Published online by Cambridge University Press