

In Epistemic Networks, Is Less Really More?

Sarita Rosenstock, Cailin O'Connor, and Justin Bruner

February 9, 2016

Abstract

We show that previous results from epistemic network models (Zollman, 2007, 2010; Kummerfeld and Zollman, 2015) showing the benefits of decreased connectivity in epistemic networks are not robust across changes in parameter values. Our findings motivate discussion about whether and how such models can inform real-world epistemic communities. As we argue, only robust results from epistemic network models should be used to generate advice for the real-world, and, in particular, decreasing connectivity is a robustly poor recommendation.

1 Introduction

Recently, philosophers of science have begun using network models to investigate the effects of communication structure on inquiry in epistemic communities.¹ These models have been taken to inform which communication structures among epistemic agents will be more or less successful in that they lead to correct beliefs about the world or facilitate consensus building in epistemic groups.

Some of the earliest, and most influential, models in this literature are presented by Zollman (2007, 2010), who argues that, under some circumstances, it will actually benefit an epistemic community to have a less connected communication network. He shows that under certain modeling conditions researchers can, by maintaining fewer lines of communication, improve the chances that as a group, they eventually come to believe true things about the world. A well-networked group of researchers, on the other hand, will tend to arrive at consensus more quickly in these models, but will be more prone to error.

This paper will highlight some of the ways Zollman's results, and related results from Kummerfeld and Zollman (2015), are sensitive to parameter choices. As we will argue, our exploration sharpens the original result. We find that in these models, less

¹We follow previous authors by taking 'epistemic community' to refer to groups of interacting researchers such as those in academia or industry, and 'network' to refer to a modeling construction in which agents are represented as points on a graph with edges denoting that the connected agents communicate with one another. For a non-exhaustive list of work in this area see Zollman (2007, 2010, 2013, 2011); Grim (2009); Kummerfeld and Zollman (2015); Holman and Bruner (2015); Mayo-Wilson et al. (2013); Alexander (2013).

connectivity improves inquiry only in a small parameter range in which learning is especially difficult: situations in which there are relatively few researchers, relatively small batches of information collected at a time, and small differences between the success of the correct and incorrect theories that the researchers are comparing. When inquiry is easier, decreased connectivity simply slows learning and provides no particular benefits.

We will argue that our results, and previous results demonstrating the sensitivity of epistemic network results to structural changes, have implications for how these models should be taken to inform real epistemic communities. In particular, we will argue that epistemic network models cannot give specific prescriptive advice to epistemic communities as to which communication structures are best for inquiry. Because results in these models are not robust across parameter or structural choices, they will yield very different advice for epistemic communities of different sorts, and for different problems these communities tackle. It is impractical, and usually impossible, however, for epistemic communities to know what sort of situation they in. Furthermore, a coarse grained assessment of the community will not do as, for these highly idealized models, the model-world match is not a close one.

As we will also argue, however, this does not imply that epistemic network models are useless. Robust phenomena across these models are much more promising vis-à-vis informing real world communities. The observation by Zollman (2010) that transient cognitive diversity improves inquiry, for example, is robust across models and seems to capture an important aspect of success in real epistemic communities. Our results suggest two more such robust phenomena. First, when inquiry is easier, network structure matter less for successful inquiry. Second, for any community, there are better ways to improve exploratory success than to decrease connectivity.

We start in section 2 by describing Zollman’s models in detail and discussing some of the ways his results have been used in subsequent literature. We highlight, in particular, work that takes his models to have prescriptive implications for epistemic communities. Section 3 presents the results of replicating Zollman’s models in a wider parameter space. In 3.1 we show that when one shifts the relative success rates of theories in Zollman’s models, his original results do not hold. In 3.2 we point out that decreasing the amount of information actors are able to gather each round increases his result and decreasing this information increases the effect. In 3.3 we show that when different assumptions are made about the size of epistemic networks, again the results break down. In section 4 we look at a similar model from Kummerfeld and Zollman (2015) and show that, here too, differences in network performance disappear under changes in parameter settings. We conclude by discussing how model robustness should be taken to inform the relevance of network epistemology models to the real world.

2 Previous Models, Results, and Impact

2.1 Models

The model used in Zollman (2007, 2010) was first introduced by Bala and Goyal (1998). These authors consider the general conditions under which a network of learning agents will adopt the same beliefs, and the conditions under which these learned beliefs will be successful ones. We will describe the exact version of the model employed by Zollman (2007), and then the version employed by Zollman (2010).

To motivate the model, consider the following situation. A group of clinicians all use drug X, with a known success rate, to treat diabetes, but a pharmaceutical company has recently released drug Y, with an unknown success rate, for the same purpose. Each clinician has some belief, based on previous experience with drug X and new data on drug Y, as to which is more successful. Each clinician begins to prescribe the drug he or she prefers to patients, all the while noting the success rates of the new drug. At the same time, clinicians communicate with colleagues about their practices and so gain information about the new drug’s success that way as well. If, as a result of this information, a clinician becomes convinced that she is doing the wrong thing, she will switch drugs. In a set-up like this, the agents will learn from their neighbors, as well as from their own choices, which action to take in the future. As such, the structure of their communication network may influence the outcome of their inquiry.

This sort of scenario can be modeled using what is called a “two-armed bandit model” as follows. Imagine a network of agents who may take one of two actions, A or B. A yields a good outcome with probability $p_A = 0.5$. B yields a good outcome with probability $p_B = 0.5 + \phi$ where $\phi \in [0, 0.5]$. However, agents do not have full information about the success of action B. They believe that it either is successful with probability $0.5 + \phi$ or with probability $0.5 - \phi$. In other words, they are unsure about whether this action is better or worse than A.

Agents in this model start with some randomly chosen belief about whether A or B is preferable (whether $p_B = 0.5 + \phi$ or $p_B = 0.5 - \phi$). This belief is modeled as a number between 0 and 1. For example, an agent might have a degree of belief of 0.8, meaning that she thinks there is a 0.8 probability that B is the better action (that $p_B = 0.5 + \phi$). Agents then repeatedly choose which action to take. In each round, agents use their degree of belief to choose an act. If it is less than 0.5, they choose A, if it is greater than 0.5, they choose B.²

Upon making a choice in each round, an agent is rewarded with a payoff selected from a Bernoulli distribution of parameters p and n . What this means is that in a single round an agent will select her preferred action n times, and each time will get a payoff of ‘1’ with probability p . For example, suppose that $p_B = 0.65$ and an agent chooses action B in a round. If $n = 100$, that agent will pick B 100 times, each time receiving ‘1’ with probability 0.65.

After acting, agents observe their outcomes and the outcomes of neighbors in their

²To be more specific, agents maximize their expected payoff for each round given their beliefs.

network. They use Bayesian updating to alter their degrees of belief about A and B based on the success rates of these actions. This sort of updating involves the application of Bayes' rule to an agent's prior degree to belief, using all the evidence gathered by an agent and her neighbors.

Upon simulating this model, agents in a network will tend to slowly converge to the same action. In other words, all agents will learn to choose A or all agents will learn to choose B. Zollman (2007) ran simulations of this sort, ending the simulations either when all agents chose A (thus ending the acquisition of new information that might lead them to choose B), or when all agents had a very high degree of belief that B was the better action (making their reversion to A vanishingly unlikely).

Zollman (2010) employs a slightly more sophisticated version of this model. In this version, agent beliefs are modeled using what is called a beta distribution. This is a distribution from $[0, 1]$ where the value of the distribution at a point represents the agent's belief that that number is the real success rate of choice B. At the beginning of a simulation, agents are given a random beta distribution with parameters α and β chosen from the interval $[0, 4]$. These parameters determine the shape of the distribution, though it is not important for our purposes to understand the details of how they do so. In each round of simulation, these agents then use their distribution to decide which action they think will be better. After acting, as before, they use Bayesian updating to alter their distributions based on their outcomes and the outcomes of their neighbors. Zollman (2010) runs these simulations for 10000 rounds before stopping them to see what agents have learned.

2.2 Zollman's Results

The central result of Zollman (2007), which is replicated in Zollman (2010), is that the network structure of such communities influences inquiry in significant ways. In particular, Zollman argues that a sparser network structure can benefit an epistemic community. More connected networks converge to uniform beliefs more quickly, but sparse networks are more likely to converge to the true belief.

Zollman (2007) achieves these results by running the models described above for agents in different network structures. First he considers three networks with three to ten agents—the cycle, the wheel, and the complete graph. Figure 1 shows these three types of networks (with eight agents in this case). He also considers every possible network formation with six agents. Figure 1 also displays two examples of size six network formations. For both sets of networks, and for both types of models, Zollman finds that more connections speed the learning of the community, but decrease the chances that the community learns to take the better action.

Why these effects? As Zollman argues, in highly connected networks, runs of unlucky results can spread quickly and convince an entire community to start taking the wrong action. In other words, the group can too quickly learn the wrong thing. In sparser networks, this is less likely to occur because agents influence each other less. It is more likely that at least a few agents maintain the better belief, and these then have a chance in the future to spread their successful beliefs to others. In a real epistemic community,




Figure 1: Various network formations. A) shows an eight agent network with complete graph, wheel, and cycle configurations respectively. B) shows two six agent networks, one more sparse, the other more connected.

this translates to pockets of interacting researchers who may maintain different sets of beliefs, but eventually share successful ones.

Zollman (2010) provides extended results using the second class of models. He shows that, in particular, it benefits communities of epistemic agents to hold a diversity of opinions for some period of time. In sparse networks, this is achieved because agents are less likely to change each others' minds quickly. It can also be achieved, he argues, by ensuring that agents begin simulations with more extreme beliefs, ones that are harder to shift.

2.3 Impact

Zollman (2007, 2010) has been widely cited both in philosophy and related disciplines. His results have been taken, in a number of cases, to suggest that a sparse network structure in epistemic communities may indeed help promote successful inquiry, and that a diversity of beliefs may be beneficial as well. In this section, we give a few examples just to demonstrate that the results of the above models have, in fact, been taken at face value by philosophers in some cases. The intent is to motivate the importance of discussing the robustness of said results, and the application of these results to real world communities.

With respect to the result that less connectivity may help a research community, Strevens (2010) writes that, "It might, for example, be preferable for scientists not to take into account too much information about their colleagues' beliefs about a problem, if a few authoritative pronouncements would stifle much-needed diversity in the range of

approaches to a problem...The question how to tune attention to authority in the short term so as to find a level of diversity that maximizes correctness in the long term has been explored with considerable insight by Zollman (2007)” (20). Wray (2013) reports that, “Zollman found that full communication in a research community is sub-optimal, as lines of communication in a group not only aid in the spread of truth, but also facilitate the spread of errors” (78).

As to the related result that diversity of beliefs among agents can promote inquiry, Douven and Kelp (2011) say that, “Zollman’s intriguing work shows that, from a socio-epistemic perspective, it may be important to maintain epistemic diversity in a community of agents, at least for a while, and that, for that reason, it is not always best if all agents have access to all information available in their community; for the same reason, a certain dogmatism on the part of the agents...may be beneficial” (277). Muldoon (2013) writes that, “Zollman showed how transient diversity in beliefs, whether fostered by limited communication or stubborn scientists, can help ensure that scientific consensus tracks the truth ” (123). And O’Connor and Bruner (2015), in motivating the importance of various sorts of diversity to epistemic communities, report that, “Zollman (2010) uses decision and game theoretic models to argue that epistemic communities that hold a variety of beliefs about the world may be more likely to eventually converge on successful theories” (3).³

To some degree at least, this work is being taken to potentially inform real epistemic networks.⁴ As we will argue in the next section, these central results of Zollman’s work depend, in some cases, on choices of model parameters.

3 Expanding the Parameter Space

In what follows, we present the results replicating Zollman’s 2007 simulations with a wider parameter space. We find that parameters for which there is a notable benefit to decreased network connectivity occupy a relatively small niche of the total space.

We ran Bala and Goyal style simulations varying p_B , n (the number of trials an agent performs each round), network configurations, and network sizes.⁵ For each parameter

³At the end of the paper we will argue that this latter result is justified in its application to real world communities. This second set of quotes is not to show that a mistake is being made, but rather just to underscore the degree to which epistemic network results are being taken seriously.

⁴Zollman’s original papers invite such an interpretation. He writes, for example, that, “...this model suggests that in some circumstances there is an unintended benefit from scientists being uninformed about experimental results in their field...When we want accuracy above all else, we should prefer communities made up of more isolated individuals” (15). We should be clear that Zollman is generally careful, in his work, not to overstate the significance of these results, but neither does he present them as how-possibly results, or something of that nature. As we will outline, he also makes clear in later papers that epistemic network models are sensitive to structural change and that, as a result, the prescription, ‘decrease connectivity for increased success of inquiry’ is not a general one.

⁵We ran simulations for a variety of parameters, with p_B ranging from 0.501 to 0.7, n ranging from 1 to 6000, and network size ranging from 4 to 100. We were constrained by computing power in some cases. For example, lower values of n , and p_B , and larger networks took longer to run. Generally, though, trends we noted held for all parameter values explored. Results reported here are winnowed from these,

setting, we ran at least 10000 trials and checked to see whether each had converged to the favorable outcome (and how quickly). We focused on two network configurations—the cycle and the complete graph. We chose these configurations as representative of highly dense and sparsely connected networks, and because both were considered in Zollman’s original paper. Where his results occur, they ought to occur for these two networks.⁶ We also ran similar simulations mimicking the models used in Zollman (2010), but for smaller sets of parameter values. The results of these simulations were similar in all cases to those we present in more detail here.

We will refer to the positive difference between successful convergence in the cycle network versus the complete network as the “Zollman effect” for the rest of the paper, as this measurement represents for these simulations the general effect noted by Zollman—that sparser networks are more reliable than well connected ones.

3.1 Success Rates

In Zollman (2007, 2010), the probability that action A is successful is always $p_A = 0.5$. The probability that action B is successful is always chosen to be $p_B = 0.501$. This means that in models from Zollman (2007), agents believe that B is either successful with probability 0.501 or 0.499. In Zollman (2010), agents’ beliefs are modeled using a more sophisticated distribution, so this is not necessary, but it is still the case in these models that $p_B = 0.501$.

Obviously, 0.5 and 0.501 are very close values. In this section, we show that in many cases as p_B increases by even a relatively small amount, the Zollman effect no longer obtains. The intuitive explanation for this is that cases with a higher p_B are cases where agents can more easily gain accurate information about the world. Trials of action A and action B are less likely to have similar results, and so agents have an easier time distinguishing them. As a result, a sparse network structure is unnecessary to help the community learn correct beliefs, but still slows learning significantly.

Figure 2 shows outcomes of our simulations when we vary p_B , but keep population size and n as in Zollman’s original models, fixed at 6 and 1000 respectively. The x-axis tracks the value of $p_B - 0.5$, the differential success rate for the better action. The y-axis shows, for the complete and cycle networks, how frequently they converge to the successful action.

As discussed, as p_B increases, the Zollman effect decreases. In other words, the difference in likelihood of successful convergence between the two network configurations decreases. This difference is effectively 0 for $p_B \geq 0.53$. For $p_B = 0.51$ there is only a difference of about 2% in the successful convergence rates.⁷ As discussed, this occurs

but the trends should be thought of as generally applicable. Recorded results and the python code used to generate them can be found here: *removed for review*

⁶It is also helpful to note that Zollman (2007) ran simulations of all networks of size six and found a monotonic relationship between sparsity and the probability of convergence to the correct theory. This suggests that for small networks in between the cycle and wheel, intermediate outcomes should be expected.

⁷Zollman (2010) says, with respect to the Zollman effect and different values of p_B , “The degree

Effect of p_B on Convergence




Figure 2: Simulation results showing the rates of convergence (*i.e.*, portion of our simulations for which the agents converged to the better action) as a function of p_B for the complete and cycle networks of size 6 with $n = 1000$.

because as p_B increases all networks are able to successfully converge to correct beliefs with very high probability, meaning that sparse networks are not particularly helpful.

3.2 Number of trials per round

Figure 2 uses Zollman’s choices of $n = 1000$ for the number of times per round that each agent tries her preferred action. In expanding the parameter space investigated for these

of difference here should not be taken too seriously; it can be altered by modifying the difference in objective probabilities of the different methodologies. However, the ordering of the graphs remains the same—the cycle is superior to the wheel which is superior to the complete graph” (17). Zollman (2007) includes the following footnote, “The results for both reliability and speed are robust for [the cycle, wheel, and complete networks] across...difference in payoff between the good and uninformative actions. Although these different modifications do effect the ultimate speed and reliability of the models, for any setting of the parameters the relationship between the three networks remains the same” (10). Given the results here we think it more accurate to say that the effect does not arise for significant portions of the relevant parameter space. Note, though, that we never see the effect reverse so it may be that it holds, but to such a slight degree that it is undetectable given the limits of our computational set up.

models we found that decreasing n actually increases the Zollman effect across models. Figure 3 shows the simulation results for a population of size 10 with p_B held fixed at 0.51, where n ranges from 10 to 6000. As this figure shows, the difference in rates of successful convergences for the cycle and complete networks is more significant for lower n .

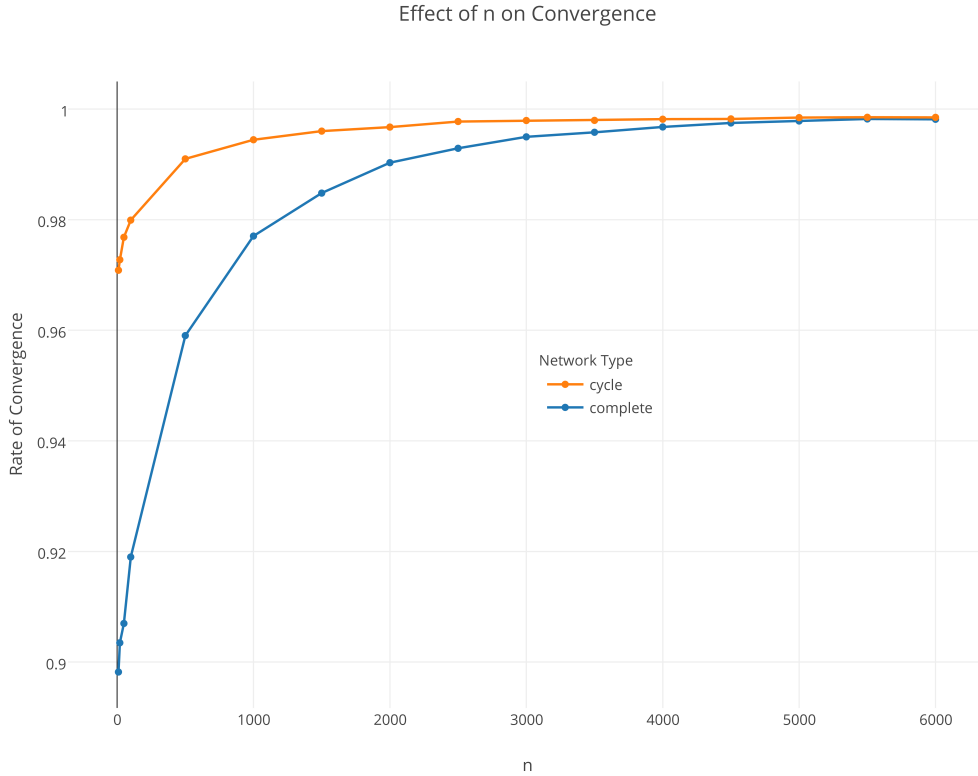


Figure 3: Simulation results showing the influence of n , or the number of trials per round, on convergence for complete graphs and cycles for a population of size 10 with $p_B = 0.51$.

In these simulations, for larger n , probabilistically it is more likely that the data obtained by each actor will reflect the true underlying probabilities of success for each action, and ditto for the combined information between actors. For smaller n , actors are more likely to receive strings of data where the less preferable action is successful, or the preferable one is unsuccessful, as a result of the law of small numbers. In a situation like this, network sparsity helps the entire community avoid these misleading strings.⁸

Even for simulations with low n , the Zollman effect disappears as p_B grows. To

⁸Note, holding fixed other parameters and varying n , actors must receive about the same numbers of data points to converge on a set of beliefs. So, for example, when $n = 20$ and actors are in a cycle

be clear, in the graphs we present we focus on areas of the parameter space where the Zollman effect occurs, but decreases with a change in parameter value. There are many areas of parameter space—high p_B , low n , and, as we shall show, large networks—where no effect is seen at all. We do not show these in the figures because there would not be much to see.

3.3 Number of Agents

As mentioned, we also looked at larger populations than Zollman considered in his original papers. In these larger populations, we found that the significance of the Zollman effect, again, decreased. Once more, we ran simulations of complete and cycle networks for various parameter values, and measured what proportion of simulations (again, out of at least 10000) converged to the better action.

Figure 4 shows results of these simulations. For this figure we look at the Zollman effect for the parameters $n = 1000$ and $p_B = 0.501$, matching Zollman’s original models. As is evident, the Zollman effect is strongest for smaller networks, but as network size increases it drops off.⁹ We find similar results for other parameter values as well.¹⁰

What is going on in this case? For these larger networks, the effects of connectivity are relatively unimportant because actors are able to gain so much information about their target inquiry from their larger cohort. Thus in almost every case, these larger networks converge to correct beliefs about the world. This ends up meaning that sparser networks do not give agents a significant advantage in terms of accuracy.

Again, sparse networks tend to learn slowly, and this is especially the case for *large* networks. Figure 5 shows the difference in speed of convergence on average for cycle vs. complete networks. As is evident the cycle is often a much slower learner, despite providing relatively little, and often nothing, in terms of improved accuracy. The larger the network, the more significant this difference. This makes sense, as in larger networks information takes longer to travel along a cycle, and the difference in connectivity between cycle and complete networks is much larger. Thus as network size increases, the potential benefit of decreased communication goes down, and the cost (i.e., speed of convergence) associated with reducing network connectivity climbs dramatically. For a population of 100 researches, the convergence rate for the complete network is 99.12%, vs. 100% for

they receive 60 data points each round. If $n = 100$ they receive 300 data points each round. If the $n = 20$ network converges in 20 rounds on average (1200 data points total for each agent), the $n = 100$ network will converge in about 12 rounds on average (1200 data points total for each agent). This supports our claim that the difference in Zollman effect across these simulations is indeed due to how often they update, and in particular the fact that in low n simulations actors update more frequently and see misleading short strings of data.

⁹For some trials the Zollman effect increases from populations of size 4 to about size 10, and then begins to decrease again. This is because the difference in connectivity between the cycle and the wheel is smaller for smaller networks (*e.g.*, the difference in number of connections each agent has is just 2 vs. 3 for population size 4, but much larger for larger populations). This means that for these very small networks differences in connectivity cannot have much of an effect.

¹⁰In the figure we show, the Zollman effect exists at all values. For many other parameter values (for instance, a small increase in p_B), it entirely disappears as the size of the network increases.

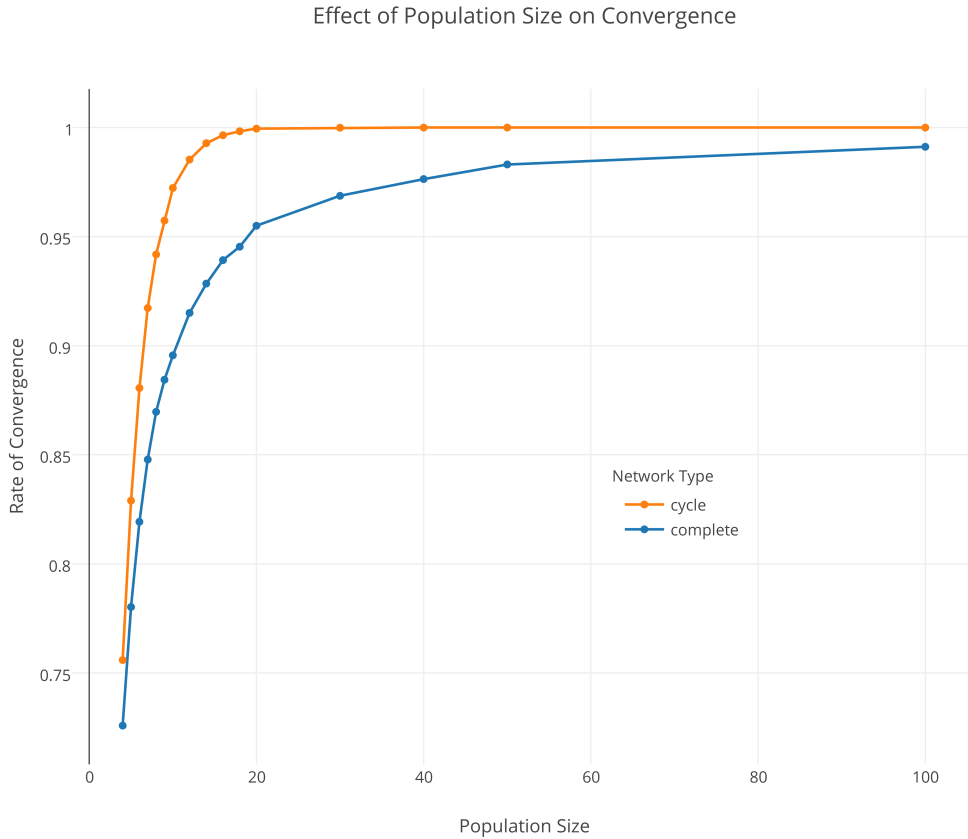


Figure 4: Rate of convergence by population size for complete graph and cycle, $p_B = 0.501$, $n = 1000$.

a cycle, while the average round at which the community achieves convergence is 20 for the complete network and 1977 for the cycle.

One argument Zollman (2007) gives is that there is a trade-off, in such networks, between learning quickly and learning well. More connected networks learn quickly, but may fail to converge to correct beliefs, while sparser networks are more accurate but slower. When the Zollman effect is non-existent, or even close to 0, this trade-off, of course, also disappears. Sparse networks still learn more slowly than densely connected ones, but there is no particular benefit from sparse network structure.

4 Exploratory Agents

In this section, we briefly turn our attention to models from Kummerfeld and Zollman (2015) who find a similar relationship between connectivity and successful inquiry in

Speed of Convergence

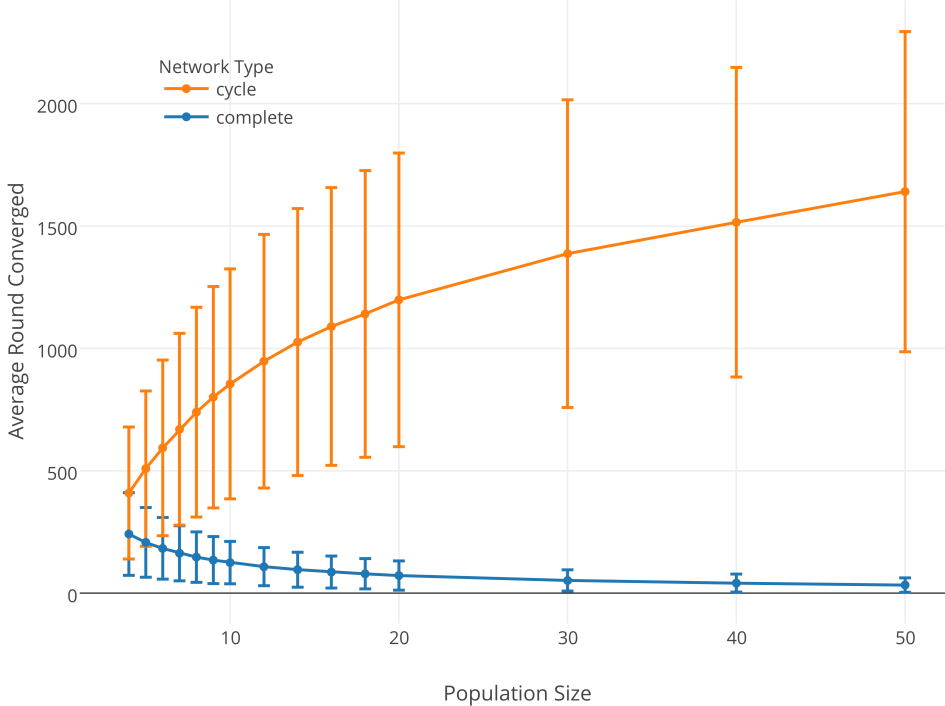


Figure 5: Speed of convergence in cycle and complete networks by population size, $p_b = 0.501$, $n = 1000$. Vertical bars show one standard deviation.

structurally different models than those we have been considering. They also show that when agents have more exploratory tendencies, complete networks are more successful than cycles (a sort of reverse Zollman effect). We will show that the effects in Kummerfeld and Zollman (2015), like those in Zollman (2007, 2010) disappear under broader parameter values.

Kummerfeld and Zollman (2015) look at models that are very similar to those discussed here. Agents are once again presented with two actions. They are part of either 8-agent complete or cycle networks. The agents in these networks, however, are not Bayesian updaters. Rather, they play a strategy called ϵ -greedy. What this means is that each agent calculates, based on her observations and those of the agents she is connected to, which action has had the highest average payoff in the past. She takes this action in the next round with probability $1 - \epsilon$, and explores, or tries the other action, with probability ϵ . Unlike the agents in the models we have considered thus far,

then, these agents do not always pick what they think is best. Instead they explore all options, with ϵ determining how likely they are to be exploratory.

The success rates of the two actions in these models also vary from those discussed to this point in the paper. The worse action always yields the same outcome of 0.¹¹ The better action is also the riskier one. It yields an outcome drawn from a normal distribution centered on a value greater than 0.

Kummerfeld and Zollman (2015) find that for low ϵ (the parameter values most like the models in Zollman (2007, 2010)), cycle networks are more successful than complete networks in that, on average over all rounds and all agents, the payoff is higher. Note that this is a different measure of success than we have been discussing, but one that also tracks whether agents are taking the better action. As ϵ increases, though, complete networks start outperforming cycles. In other words, as agents explore more on their own, connectivity starts to be beneficial to inquiry.

We looked at an expanded range of parameter values for these models and, again, found that network effects disappeared for wider parameter values. Kummerfeld and Zollman (2015) looked at models where the better action resulted in a payoff drawn from a normal distribution centered on 1 with a standard deviation of 9. We explored a much less extensive parameter space for these models than the Bala and Goyal style ones, but we looked at models where this distribution instead had a mean = 1, 2, 3, and 4, and standard deviation = 3 and 9. We considered a version with more exploratory agents, $\epsilon = 0.35$, where the complete network was better, and less exploratory agents $\epsilon = 0.05$, where the cycle network was more successful instead.¹²

We found that as the mean of the payoff distribution increased, network effects on performance decreased both for the low and high ϵ . We also found that decreasing the standard deviation of the distribution decreased network effects. Figure 6 shows these results. The x-axis represents the mean of the payoff distribution from 1 to 4. The y-axis tracks the difference in success between the cycle and complete network for each parameter value. As the mean increases, the difference between the two networks decreases for all parameter values. Also, as the standard deviation decreases, the difference in success become less.¹³

Once again, we find that in the larger parameter space of this network epistemology model, the effects of connectivity disappear over many values. In this case, we see the disappearance both of a Zollman effect, and a sort of reverse Zollman effect where connectivity improves inquiry. Notice that in these simulations again, when agents have an easier time determining the better action—when the mean is higher and the standard

¹¹The value 0 should not be taken to mean this is a useless action, but rather is a baseline of comparison for the other arm.

¹²We looked at networks of size 8. Each simulation was run for 1000 rounds. In each round, actors performed the action they thought best 50 times, with probability $1 - \epsilon$. We ran 100 simulations at each parameter value. The reported success rates are the average payoffs for all actors over all rounds of simulation.

¹³We confirmed this for other parameter values with networks of size 7 using fewer rounds. This was to decrease computational requirements. For these networks again, increasing the mean and decreasing the standard deviation of the payoff distribution decreased network effects for both low and high ϵ .




Figure 6: Difference in performance between cycle and complete networks as the mean of the payoff distribution for the better action increases.

deviation lower—network structure is less important.

5 Discussion

In this section, we first pull out some relevant lessons from our extended analysis of the models in Zollman (2007, 2010) and Kummerfeld and Zollman (2015). We then turn to the implications of this work, and previous work, on the usefulness of epistemic network models to informing real world epistemic communities.

5.1 When Does Network Structure Matter in the Models?

In Bala and Goyal style models of epistemic networks, less connectivity *can* improve the accuracy of learning, but this only happens for certain areas of parameter space. In particular, the effect occurs under circumstances where learning is more ‘difficult’ for the agents in that the community has more trouble distinguishing between two alternative theories.

The learning situation can be more difficult in this way when the parameters have the following features:

1. The two actions that agents may take are more similar in terms of success rates (low p_B);

2. The population size is smaller; and
3. The amount of data collected each round is smaller (low n).

In the first case, greater amounts of data are necessary to accurately determine which action is the preferred one. In the second and third cases, agents receive less data each round. In all three cases, (relative) sparsity of data allows for misleading strings of information to sway epistemic communities to incorrect beliefs, meaning that there is an opportunity for network structure to influence inquiry.

To illustrate this claim consider figure 7. For each set of parameters— n , network size, and p_B —we average convergence times for the cycle and wheel network using our data from Section 3. This average gives us a proxy for how difficult it is for networks of agents to converge under each set of parameter values. We then compare this measure of speed to the strength of the Zollman effect. The data is divided into smaller networks (6, 10, and 20) and larger ones (50, 100) since there is a significant difference in the strength of the effect in these two sets of networks. Note that the y-axis is on a logarithmic scale to make the trend more visible. As is evident, there is a clear correlation between networks that take more time to converge and those where connectivity matters. This supports the claim above that network structure matters more when inquiry is trickier.

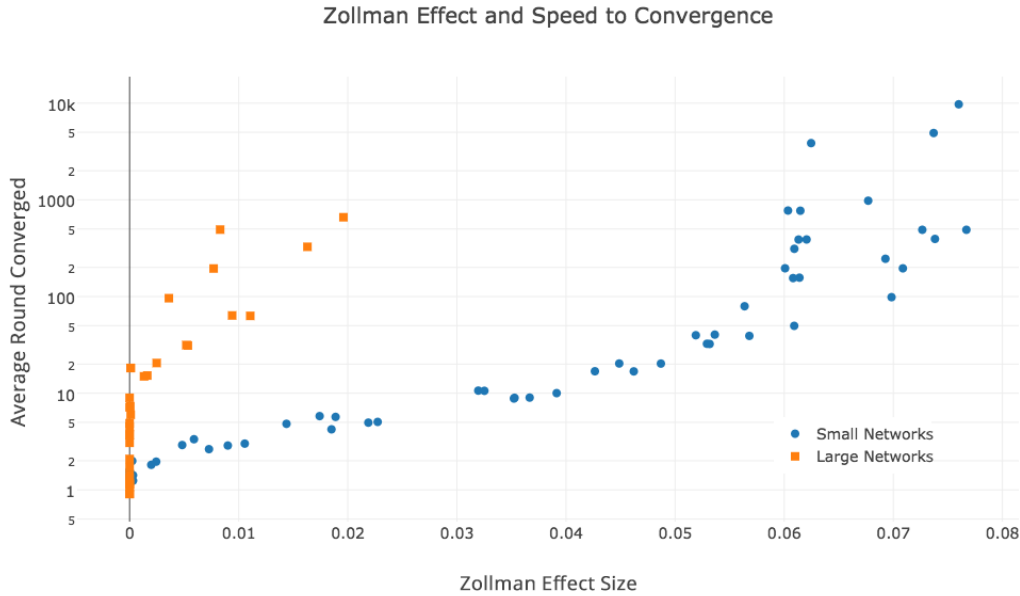


Figure 7: Average speed of convergence over cycle *and* complete networks as a function of Zollman effect size.

There is something unintuitive about this observation. When good information is harder to come by, this is exactly the situation in which, for these models, it is useful

to decrease the amount of information flowing between agents at each time step. The way to think about this is to observe that sparsity in epistemic networks can provide the very benefit that Zollman outlined (helping groups of agents to avoid preemptively converging to incorrect beliefs about the world), but it will only provide it when agents are already in a more difficult situation for inquiry. When agents have enough good data, decreasing connectivity only slows learning without providing any benefit.

In Kummerfeld and Zollman (2015) type models, again, network effects only show up when it is more difficult for agents to distinguish between alternative theories. When the mean of the payoff distribution for the more successful action is similar to the payoff for the unsuccessful action, both a Zollman effect and a reverse effect are seen for different levels of exploration. As this mean becomes larger, i.e., the better action is more obviously better, the different networks become equally successful. When the standard deviation of the payoff distribution is large—when it is harder to determine that the better action is, in fact, better—again both sorts of network effects are seen. And as standard deviation decreases, networks matter less.

5.2 Real World Implications

Now we turn to the following question: what are the implications of our results on the use of epistemic network models to inform the real world? To begin with, we should distinguish between different sorts of purposes that models of epistemic inquiry (and models more generally) can serve. Such models can be employed to generate explanation of processes in real epistemic communities. Alternatively, they can be used to provide prediction and prescription for epistemic communities regarding questions like, ‘what communication networks support successful inquiry?’ We will address the potential for epistemic network models to play this latter sort of role.

What sorts of real world recommendations do network epistemology models license? Obviously it would be a mistake to go from Zollman’s original work to the conclusion that real epistemic communities should generally decrease connectivity between agents, or from Kummerfeld and Zollman (2015) to the conclusion that epistemic communities should increase connectivity.¹⁴ At very least, our results indicate that such measures should only be helpful for communities confronted by a more difficult situation for inquiry in the ways outlined above. Yet even this is too strong. Before delving into the details, it is helpful to first discuss a distinction from Weisberg and Reisman (2008). In the course of this paper we have demonstrated that certain results in Zollman’s models lack what these authors call *parameter robustness*, or the invariance of an effect across possible parameter values. They differentiate this from *structural robustness*, or the invariance

¹⁴We should mention that previous authors would agree with this. In a survey paper discussing results from network epistemology models Zollman (2013) points out that, “A remarkable amount depends on the underlying learning problem. In some situations (like the pooling examples from Section 3.1), more communication is better. In other situations (like the information transmission problems from Section 2 or bandit problems from Section 4.2) less communication is better. Sometimes increased communication is harmful because it comes at a cost and fails to improve the epistemic performance of the group” (25).

of an effect across varying structural assumptions in a model.¹⁵

Does the Zollman effect show structural robustness? Grim (2009) finds something similar to the Zollman effect for structurally different epistemic networks (in particular, individuals aren't presented with just two actions but instead are placed on an 'epistemic landscape'). As discussed, Kummerfeld and Zollman (2015) find that under some structural changes—agents who are not Bayesian learners, and a different payoff structure—the Zollman effect is replicated. Under other structural changes, however—the introduction of naturally exploratory agents into epistemic networks—the Zollman effect reverses. Holman and Bruner (2015) look at Bala and Goyal style networks where some agents are not motivated by epistemic concerns. These 'biased' agents simply attempt to sway the scientific community to their views. Again, in these models connected, rather than less connected, networks are more successful. To summarize, the Zollman effect has some structural robustness, but other similar epistemic network models with structural tweaks find the opposite effect—that more connected networks are more successful.

This structural non-robustness means that even for situations in which network structure does matter, further details regarding the members of the community will dictate what the optimal network structure is. In other words, if these models are to be used to recommend specific network structures for real world epistemic communities, members of such communities must identify what sort of situation they are themselves in.¹⁶

But it is extremely difficult to know the learning situation of a community or the underlying motivation of its members, especially before inquiry begins. To make this clear, briefly consider the slew of questions regarding the structure of the community which must be addressed to determine whether a less connected network would actually improve the success of inquiry. How much data will our community be able to generate on a particular problem, and in what sized chunks will this data be generated? How many researchers will tackle a particular problem in the end? Is one theory going to turn out to be clearly better than another theory? Are there biased peers in the community? Are we exploratory agents, who will keep after a possibility even if we think it might be wrong, or not? It is impractical and often impossible to determine such details. One might respond by coarse-graining. Surely we can determine the approximate size of the community or the approximate size of the trials that can be run? Note, however, that one will not be able to determine ahead of time whether a community harbors biased agents. Furthermore, the relative success of theories to be tested (p vs. p_B) is precisely what the community is attempting to measure. It is not possible to determine ahead of time. Yet since the Zollman effect displays neither structural nor parameter robustness, all these questions, including the impossible ones, must be dealt with to determine whether epistemic network models recommend a sparser or denser network.¹⁷

¹⁵The distinction between parameter and structural robustness will not always be a useful one. The same change to a model might be considered a structural change *or* a parametric change depending on the focus of inquiry. Thanks to Louis Narens for this point.

¹⁶As Zollman (2013) concludes after reviewing the effect of network structure in several different epistemic situations.

¹⁷Along these lines, it may be easier to use these models for post-hoc explanation because it retrospect

One might propose that we can have probabilistic beliefs about what the features of our epistemic communities are. So perhaps we can use these in a decision calculus that will output an ideal network structure. This suggestion runs us into a further problem however. Since epistemic network models are highly idealized, it is not the case that parameter values or features of the model neatly correspond to real world values and features. For this reason, even if agents in a real epistemic network have beliefs about the features of their epistemic community, they will not be able to say which measured variables in the world correspond to which parameters in the model. And thus they will not be able to use these beliefs to pick a model.

We should be clear that when it comes to prescriptive advice, we do think there is a role for network epistemology models to play. There are some results from these models that are robust across parameter and structural choices. For these results, epistemic communities need not worry about how the details of their situation matches the models, and so the problems just described dissolve. Robustness can increase our confidence that the models tell us something about particular real world communities because we expect that all such communities (or many) fall under the purview of these results.¹⁸

As we have mentioned, one such result, discussed at length in Zollman (2010), is that transient cognitive diversity in epistemic communities improves inquiry. Transient diversity here refers to the idea that for at least some period of time, agents disagree about what theory is the correct one. This sort of diversity prevents epistemic communities from, incorrectly, rejecting good theories. In our models, as in previous ones, at least some level of transient diversity is necessary to ensure that epistemic communities arrive at the right choice.¹⁹ A second such result is the one mentioned in 5.1—that when inquiry is easy for various reasons, network structure is unimportant. This result can be thought of as modifying the one just mentioned. When inquiry is easier, transient diversity, as ensured by network structure, will be less necessary to ensure good results.

One last robust result from our exploration here is that, across models, there is always a better way to improve the success of inquiry than network sparsity. As we have shown, raising the population size and/or the number of trials per round (n) can achieve the same increase in rate of convergence to the better action without such a high cost in number of rounds required. And if this is not possible—if there is an external restriction on the amount of data that can be collected—then this is exactly the sort of situation in which the community cannot afford to sacrifice the extra data they would get by observing all of their peers' results, since good data is hard to come by. Instead, the most reasonable recommendation is to change the feature of the model that all agents

it will be easier to determine relevant facts about an epistemic community. Even in cases of post-hoc explanation, however, model fitting is tricky with highly idealized models as we outline below.

¹⁸Unsurprisingly, philosophers of science disagree about the role of robustness analysis in modeling. Some maintain that robust results in the sense just described are better supported (Levins, 1966; Wimsatt, 2012; Weisberg, 2006; Heesen et al., 2014). Others hold that this sort of robustness analysis, in fact, cannot support modeling results (Orzack and Sober, 1993; Odenbaugh and Alexandrova, 2011; Cartwright, 1991). We side with the former.

¹⁹One might point out that while temporary disagreement is often beneficial, the optimal amount of, and even kind of, transient diversity needed hinges on particular details of the community in question. In some cases, little, if any, transient diversity is required to ensure the group is likely to succeed.

always and only perform the action that they currently believe to be preferable in favor of determining which action is in fact better—potentially sacrificing short-term in favor of long-term utility maximization.

References

- J McKenzie Alexander. Preferential attachment and the search for successful theories. *Philosophy of Science*, 80(5):769–782, 2013.
- Venkatesh Bala and Sanjeev Goyal. Learning from neighbours. *The review of economic studies*, 65(3):595–621, 1998.
- Nancy Cartwright. Replicability, reproducibility, and robustness: Comments on harry collins. *History of Political Economy*, 23(1):143–155, 1991.
- Igor Douven and Christoph Kelp. Truth approximation, social epistemology, and opinion dynamics. *Erkenntnis*, 75(2):271–283, 2011.
- Patrick Grim. Threshold phenomena in epistemic networks. In *AAAI Fall Symposium: Complex Adaptive Systems and the Threshold Effect*, 2009.
- Remco Heesen, Liam Kofi Bright, and Andrew Zucker. Vindicating methodological triangulation. 2014.
- Bennett Holman and Justin P. Bruner. The problem of intransigently biased agents. *Philosophy of Science*, 82(5):956–968, 2015.
- Erich Kummerfeld and Kevin J. S. Zollman. Conservatism and the scientific state of nature. *British Journal for the Philosophy of Science*, Forthcoming, 2015. doi: 10.1093/bjps/axv013.
- Richard Levins. The strategy of model building in population biology. *American scientist*, pages 421–431, 1966.
- Conor Mayo-Wilson, Kevin Zollman, and David Danks. Wisdom of crowds versus group-think: learning in groups and in isolation. *International Journal of Game Theory*, 42(3):695–723, 2013.
- Ryan Muldoon. Diversity and the division of cognitive labor. *Philosophy Compass*, 8(2):117–125, 2013.
- Cailin O’Connor and Justin P Bruner. Dynamics and diversity in epistemic communities. Working Paper, 2015.
- Jay Odenbaugh and Anna Alexandrova. Buyer beware: robustness analyses in economics and biology. *Biology & Philosophy*, 26(5):757–771, 2011.

- Steven Hecht Orzack and Elliott Sober. A critical assessment of levins's the strategy of model building in population biology (1966). *Quarterly Review of Biology*, pages 533–546, 1993.
- Michael Strevens. Reconsidering authority scientific expertise, bounded rationality, and epistemic backtracking. 2010.
- Michael Weisberg. Robustness analysis. *Philosophy of Science*, 73(5):730–742, 2006.
- Michael Weisberg and Kenneth Reisman. The robust volterra principle*. *Philosophy of science*, 75(1):106–131, 2008.
- William C Wimsatt. Robustness, reliability, and overdetermination (1981). In *Characterizing the Robustness of Science*, pages 61–87. Springer, 2012.
- K Brad Wray. The future of the structure of scientific revolutions. *Topoi*, 32(1):75–79, 2013.
- Kevin JS Zollman. The communication structure of epistemic communities. *Philosophy of Science*, 74(5):574–587, 2007.
- Kevin JS Zollman. The epistemic benefit of transient diversity. *Erkenntnis*, 72(1):17–35, 2010.
- Kevin JS Zollman. Social network structure and the achievement of consensus. *Politics, Philosophy & Economics*, page 1470594X11416766, 2011.
- Kevin JS Zollman. Network epistemology: Communication in epistemic communities. *Philosophy Compass*, 8(1):15–27, 2013.