

The Communication Structure of Epistemic Communities  
Kevin J.S. Zollman

## 1 Introduction

Increasingly epistemologists have become interested in the relationship between social influences and proper epistemic behavior. The analysis of this set of issues comes in one of two forms. One form is to consider the proper response for epistemic agents when faced with evidence that comes via another person (or persons). This type of analysis remains focused on the traditional epistemic problems of individual belief formation and revision, but incorporates appropriate responses to data of a certain kind.

Another approach focuses more on the structure of epistemic communities. This second type asks, given certain assumptions about the individuals in communities, what sort of community structures best serve the epistemic aim of that community? Alvin Goldman (2009) calls epistemology of this latter sort “systems-oriented social epistemology”. As an example of this sort of epistemology, Philip Kitcher (1990, 1993) and Michael Strevens (2003a, 2002b) have recently looked at the impact that different methods for assigning credit have on communities of scientists. They conclude that our current method of assigning credit is best for achieving the desired results of science.

Here we will be interested in one among many potentially interesting features of communities, namely, the structure of communication. Specifically we will ask: what is the best way for information to be transmitted? In order to analyze this problem we will look at the prime example of an epistemic community, science. In order to do this, we will use a model first suggested by two economists, Venkatesh Bala and Sanjeev Goyal (1998). The surprising result of this analysis is that in many cases a community that withholds information from its members is more reliable than one that allows for fully informed individuals. One might expect that reducing information to scientists would also have the effect of making their convergence to the truth much slower, and our model confirms this suspicion. The model suggests that there is a robust trade-off between speed and reliability that may be impossible to overcome.

After presenting the model in Section 2, the results from a computer simulation study of the model are presented in Section 3. Section 4 discusses the limitations of the model as a model of science, and Section 5 concludes by comparing the results of this model with another problem discussed by Kitcher and Strevens.

## 2 The Model

Consider the following stylized circumstance. There are four medical researchers working on a particular disease. They are confronted with a new treatment method that might be better or worse than the current, well-understood, method of treatment. Work on the new treatment will help to determine whether it is superior. Since the old treatment is well understood, experimental work on it will not result in any new information about its probability of success; scientists' efforts will only refine delivery methods or reduce harmful side-effects. Suppose our scientists, labeled  $A$ ,  $B$ ,  $C$ , and  $D$ , respectively assign the following probabilities to the superiority of the new treatment: 0.33, 0.49, 0.51, and 0.66. Then each pursues the treatment method that she thinks best. Two scientists,  $C$  and  $D$ , pursue the new treatment option and the other two,  $A$  and  $B$ , pursue the old treatment option. Suppose, further that the new treatment is in fact better than the old but, as is perfectly possible,  $C$  and  $D$ 's experiments both suggest slightly against it.<sup>1</sup> After meeting and reporting their results to each other all the scientists might now judge it to be more likely that the old treatment is superior.<sup>2</sup> As a result, none of them will experimentally pursue the new treatment; we have lost a more beneficial treatment forever.

This circumstance arises for two reasons. First, scientists in our example must pursue evidence, they are not passive observers. Second, they already have a good understanding of the old treatment and further study of it will not help them to conclude anything about the new treatment.<sup>3</sup>

Even given this structure, the availability of the evidence contributes to the abandonment of the superior theory. Had  $D$  not been aware of  $C$ 's result, she would still have believed in the superiority of the new treatment.<sup>4</sup> As a result, had she been unaware of  $C$ 's results, she would have performed a second round of experiments, which would offer the opportunity to correct the experimental error and thereby to find the truth. In this toy example, it seems that the wide

---

1 Specifically suppose that all agree on these conditional probabilities:

$$\begin{aligned} P(\text{The result of } C\text{'s experiment} \mid \text{New method is better}) &= 0.4 \\ P(\text{The result of } D\text{'s experiment} \mid \text{New method is better}) &= 0.4 \\ P(\text{The result of } C\text{'s experiment} \mid \text{New method is worse}) &= 0.6 \\ P(\text{The result of } D\text{'s experiment} \mid \text{New method is worse}) &= 0.6 \end{aligned}$$

- 2 Using the numbers above,  $A$ ,  $B$ ,  $C$ , and  $D$  would now assess the probability of the new theory being better as 0.1796, 0.2992, 0.3163, and 0.4632 respectively. This outcome is far from extraordinary. Given that the new methodology is better and the experimental outcomes are independent (conditioned on the new methodology being superior), the probability of getting this result is 0.16.
- 3 Had the scientists been passive observers, their beliefs would not have influenced the type of information they received. In that case, information about either treatment might still arrive despite the fact that the theory has been abandoned. Additionally, had experiments on the old theory been informative about the effectiveness of the new theory, the fact that everyone pursues the old theory does not preclude them from learning about the new theory.
- 4 If  $D$  had only been aware of her own negative results, but not the results of  $C$ , her posterior belief in the superiority of the new treatment would have been 0.5621.

availability of experimental results was detrimental to the group's learning. Of course no general lesson can be drawn from this example. It is not offered as a general model for all scientific practice but is instead provided as a generalization of a learning situation that some scientists unquestionably face.

Two economists, Bala and Goyal (1998) present a very general model that can be applied to circumstances like the one faced by the medical researchers. Stated formally, in this model, there are two states of the world  $\phi_1$  and  $\phi_2$  and two potential experimental options to pursue  $A_1$  and  $A_2$ . Option  $A_1$  has the same expected return in both states while  $A_2$ 's is lower in  $\phi_1$  and higher in  $\phi_2$ . The return from choosing an option represents the degree to which a particular experiment succeeds – a higher payoff represents a larger experimental success. Agents are aware of the expected payoff in both states, but are unaware of which state obtains. Agents have beliefs about the state of the world and in each period pursue the option that has the highest expected utility given their beliefs. They receive a payoff from their actions that is independently drawn for each player from a common distribution with the appropriate mean. Each agent observes the outcome of his choice and the outcomes of *some* others, and then updates his beliefs about the state of the world based on simple Bayesian reasoning.<sup>5</sup>

This model has multiple interpretations, but one of them is analogous to the circumstance discussed above. The agents are scientists and their action is choosing which method to pursue.  $\phi_1$  represents the state where the current method is better and  $\phi_2$  is the state where the new method is better. Bala and Goyal endeavor to discover under what conditions groups will converge to taking the best action in a given state. They consider two different restrictions: restrictions on priors and restrictions on information about experimental outcomes.

The second suggestion, limiting information about outcomes, will be our primary focus here. This restriction is achieved by limiting those other agents a given individual can “see,” and thus restricting the information on which an agent can update. They do this by placing an agent on a graph and allowing her to see only those agents with whom she is directly connected.

Bala and Goyal consider agents arranged on a line where each agent can see only those agents to the immediate left and right of them. If there are an infinite number of agents, convergence in this model is guaranteed so long as the agents' priors obey some mild assumptions. Bala and Goyal also consider adding a special group of individuals to this model, a “royal family.” The members of the royal family are connected to every individual in the model. If we now consider this new collection of agents, there is positive probability that the group will converge to the worse option!

---

5 “Simple” here means that the agent only updates her belief using the evidence from the other's experiment. She does not conditionalize on the fact that her counterpart performed a particular experiment (from which she might infer the results of others).

This is a remarkable result, because it contradicts a basic intuition about science: that access to more data is always better.<sup>6</sup> In this case, it is not.

The reason for this result is interesting. In the single line case the probability that everyone receives misleading results becomes vanishingly small as the population grows to infinity. However, in the population with the royal family, this probability no longer vanishes. Negative results obtained by the royal family infect the entire network and mislead every individual. Once the entire population performs act  $A_1$ , they can no longer distinguish between the good and bad states because this action has the same expected payoff in both  $\phi_1$  and  $\phi_2$ . As a result a population composed entirely of  $A_1$  players will never escape.

One might worry about Bala and Goyal's results since they depend so critically on the infinite size of the population. For finite populations, there exists a positive probability that *any* population will not learn the correct action. One might wonder, in these cases how much influence the "royal family" would have. Furthermore, it is unclear what moral we ought to draw from these results – many things are different in the two different models. In addition to increased connectivity, there is also unequal distribution of connections. If we are interested in evaluating the performance of actual institutions it is unclear which features we should seek out. Through computer simulations, we will endeavor to discover the influence that network structure has on reliable learning in finite populations and also to develop more detailed results regarding the relationship between network structure and success.

### 3 Finite Populations

#### 3.1 The "Royal Family" Effect

[Figure 1 here]

**Figure 1:** A 10-person cycle, wheel, and complete graph

To begin, we will look at three graphs known as the cycle, the wheel, and the complete graph (pictured in Figure 1) and compare their convergence properties. The cycle is a finite analogy to

---

<sup>6</sup> Ellison and Fudenberg (1995) present a different model that comes to the same conclusions. In their model, the interaction structure is not fixed; individuals take a different random sample of fixed size in each time period. Because the individuals in their model have much shorter memories, it seems less appropriate for modeling scientific behavior (an application which they do not consider). A similar conclusion can be found even for individual learning in the work of Herron, Seidenfeld and Wasserman (1997). This work presents a rather different learning situation and will not be discussed in detail here

Bala and Goyal's line. Here agents are arranged on a circle and only connected with those on either side of them. The wheel is a cycle but one of the agents – Bala and Goyal's royal family – is connected to everyone else. The last network is one where everyone is connected to everyone.

We will, unbeknownst to our agents, make the world be  $\phi_2$ , where the new methodology is better. We will then assign our agents random beliefs uniformly drawn from the interior of the probability space and allow them to pursue the action they think best. They will then receive some return (a “payoff”) that is randomly drawn from a distribution for that action. The agents will then update their beliefs about the state of the world based on their results and the results of those to whom they are connected. A population of agents is considered to be finished learning if one of two conditions are met. First, a population has finished learning if every agent takes action  $A_1$ , in this case no new information can arrive that will convince our agents to change strategies. (Remember that the payoff for action  $A_1$  is the same in both states, so it is uninformative.) Alternatively the network has finished learning if every agent comes to believe that they are in  $\phi_2$  with probability greater than 0.9999. Although it is possible that some unfortunate sequence of results could drag these agents away, it is unlikely enough to be ignored.

[Figure 2 here]

**Figure 2:** Learning results of computer simulations for the cycle, wheel, and complete graphs

The results of a computer simulation are presented in Figures 2 and 3. In Figure 2, the x-axis represents the total number of agents and the y-axis represents the proportion of 10,000 runs that reached the correct beliefs.<sup>7</sup> The absolute probabilities should not be taken too seriously as they can be manipulated by altering the expected payoffs for  $A_1$  and  $A_2$ . On the other hand, the relative fact is very interesting. First, we have demonstrated that Bala and Goyal's results hold in at least some finite populations. In all the sizes studied the cycle does better than the wheel. Second, we have shown that both of these do better than the complete graph where each agent is informed of everyone else's results.

[Figure 3 here]

**Figure 3:** Speed results of computer simulation for the cycle, wheel, and complete graphs

---

<sup>7</sup> Although it is possible for a population to continue unfinished indefinitely, no population failed to converge.

This demonstrates a rather counterintuitive result, that communities made up of less informed scientists might well be more reliable indicators of the truth than communities that are more connected. This also suggests that it is not the unequal connectivity of the “royal family” that is the culprit in these results. The harm done by the individual at the center cannot be simply overcome by removing their centrality.

There is a benefit to complete networks, however; they are much faster. Figure 3 shows the average number of generations it takes to reach the extreme beliefs that constituted successful learning among those networks that did reach those beliefs. Here we see that the average number of experimental iterations to success is much lower for the complete network than for the cycle, and the wheel lies in between. This suggests that, once networks get large enough, a sacrifice of some small amount of accuracy for the gain of substantial speed might be possible.<sup>8</sup>

[...]

## 3.2 Connectivity and Success

[...]

Why is it that less connected networks, like the cycle and the wheel, are superior to more connected ones, like the complete graph? It appears that sparsely connected networks have a much higher “inertia.” This inertia takes two forms. First, a less connected network experiences fewer widespread changes in strategy on a given round than a highly connected network. The average number of people who change their strategies after the  $A_2$  players receive an unlikely low return is four times higher in a highly connected network than a less connected network. Second, unconnected networks are less likely to occupy precarious positions than connected ones. Conditioning on the network having only one  $A_2$  player, a highly connected network is almost three times as likely to have no individuals playing  $A_2$  on the next round. Since there is only one new piece of evidence in both cases, the difference between the two networks is the result of individuals having less extreme beliefs (i.e., closer to 0.5) in the connected network. Since all networks have the same expected initial beliefs, this must be the result of the information received by the agent.<sup>9</sup>

---

8 The results for both reliability and speed are robust for these three networks across modifications of both the number of strategies (and thus states) and the 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.

9 The statistics reported here are comparing 100 runs of a complete six person network with the most reliable six person network pictured in Figure 4.

Both of these results suggest that unconnected networks are more robust to the occasional string of bad results than the connected network because those strings are contained in a small region rather than spread to everyone in the network. This allows the small networks to maintain some diversity in behaviors that can result in the better action ultimately winning out if more accurate information is forthcoming. This also explains why we observed the stark difference in speeds for the cycle and complete networks in the previous section. When bad information is contained so too is good information. In fact, we find that this trade-off is largely robust across networks.

[Figure 4 here]

**Figure 4:** The five most accurate (top) and five fastest (bottom) networks

An inspection of the five most reliable and five fastest networks suggests that the features of a network that make it fast and those that make it accurate are very different (see Figure 4). Four of the five most reliable graphs are minimally connected – i.e., one cannot remove any edge without essentially making two completely separate graphs. Conversely, the five fastest graphs are highly connected, two of them are complete graphs, and the remaining ones are one, two, and three edges removed from complete graphs. Figure 5 compares the average time to success and probability of success for all networks of size six. Here we find that there is a relationship between the accuracy of a network and its speed. In fact, this graph shows that sometimes a small increase in probability can result in a substantial increase in time to success.

[Figure 5 here]

**Figure 5:** Speed versus accuracy for networks of size six

This confirms the trade-off suggested before, in order to gain the reliability that limiting information provides, one must sacrifice other benefits, in this case, speed. In fact, the trade-off is even stronger than suggested here. These results are only for cases where we specify that the new method is better. When the uninformative action is better convergence is guaranteed but the connectedness of the graph determines its speed.

[...]

Ultimately, there is no right answer to the question of whether speed or reliability is more important – it will depend on the circumstance. Although a small decrease in reliability can mean a relatively large increase in speed, in some cases such sacrifices may not be worth making. If it is critical that we get the right result no matter how long it takes we would prefer groups where information is limited (without making the network disconnected). On the other hand, if speed is important and correct results are not as critical perhaps a more connected network is desired. It is not the intention of this study to provide unequivocal answers to these questions, but rather to demonstrate that such trade-offs do exist and that one can achieve increased reliability by limiting information.

## 4 The Right Model

There are four assumptions that underlie this model that might cause some concern. They are:

1. The learning in this model is governed by the observation of payoffs.
2. There is an uninformative action whose expected payoff is well known by all actors.
3. The informative action can take on one of very few expected payoffs and the possibilities are known by all actors.
4. Individuals are *myopic* – on each round they take the action they think is currently best.

The first assumption is of little concern. Here we use payoffs to symbolize experimental outcomes. Payoffs that are closer to the mean are more likely, which corresponds to experimental outcomes that are more likely on a given theory. The payoffs are arranged so that an individual who maximizes her expected payoff pursues the theory that she thinks is most likely to be true. This fact allows this model to be applied to learning situations where individuals are interested in finding the most effective theory (however effectiveness is defined) and also to situations where individuals are interested in finding the true theory. In either case the individuals behave identically.<sup>10</sup>

The second and third assumptions are less innocuous. In another paper, I show that similar results obtain in a model where the second and third assumptions are relaxed. Even in cases where both actions are unknown and can take a variety of values, information is harmful (Zollman 2009).

The final assumption is likely to make a difference, although as far as I am aware, there has not

---

<sup>10</sup> This is not to say that true theories always have higher payoffs. Instead, this model is so general as to apply to either circumstance.



been specific investigation of variations along this line. In the model presented here the individual scientists are choosing the option that currently looks best; they are not considering the value that might be gained from pursuing the apparently suboptimal action in order to gain additional information. If an individual thinks we are probably in state  $\phi_1$  but is unsure, she might want to pursue action  $A_2$  to ensure that her current belief is indeed correct. We exclude this possibility, but why should we?

First, I think the assumption of myopia more closely accords with how individual scientists choose methodologies to pursue. In situations of this sort, calculating when it is best to take a sub-optimal action for the benefit of additional information can be *very* complex (cf. Berry and Fristedt 1985). The assumption of myopia corresponds to individuals who are willing to do some computation to determine the best course of action, but are not willing to engage in very complex calculations to do so. Second, myopic behavior is optimal when one cares significantly more about the current payoff and less about future payoffs. Information is only valuable when it will be put to future use. Since scientists are often rewarded for current successes (whether it be via tenure, promotion, grants, or awards), it is likely that scientists are at least close to myopic. Finally, even if this assumption is unreasonable, it represents an interesting starting point from which we can gauge the effect of making scientists less and less myopic.

These assumptions rest on a particular way of pursuing Systems-Oriented Social Epistemology. One type of social epistemology designs “scientific utopias” where everything – from social structures to individual behaviors – is in perfect harmony. The project of this paper instead takes its cue from Rousseau, who described his political philosophy as “taking men such as they are, and laws such as they may be made” (Rousseau 1791 [1954]). We are interested in knowing what institutions make the best out of potentially imperfect individuals, and so we will fix individuals’ behavior in a reasonable way and compare how these individuals do when placed in different social circumstances.

One might wonder if any situation faced by scientists actually fits this model. First, I believe this model very closely mimics Larry Laudan’s (1996) model of theory selection. Laudan suggests that theory choice is a problem of maximizing expected return. We ought to choose the theory that provides the largest expected problem solving ability. Since we have often pursued a particular project for an extended time before being confronted with a serious contender, we will have a very good estimate of its expected utility. However, we will be less sure about the new contender, but we could not learn without giving it a try.

Even beyond Laudan, there may be particular scientific problems that fit this model. Bala and Goyal compare their model to crop adoption in Africa. There, a new seed is introduced and farmers must decide whether to switch from their current crop (whose yield is well known) to

another crop (whose yield is not).

Scientists often must choose between different methodologies in approaching a particular problem and different methods have different intrinsic probabilities of succeeding.<sup>11</sup> In previous work, I connected this model of sequential decision making to a case of medical research (Zollman 2009). For some time there were two prevailing theories about the cause of Peptic Ulcer Disease. On one theory (the hypoacidity theory) peptic ulcers were caused by excess acid production in the stomach and treatment would involve finding various ways to reduce this acid. The other theory (the bacterial theory) claimed that peptic ulcers are caused by a bacterium that must be eradicated in order to cure the ulcers. These two theories competed for some time before the hypoacidity theory won, to eventually be supplanted by the bacterial theory which was resurrected almost fifty years after its initial abandonment. Scientists researching potential treatments for peptic ulcers chose experimental treatments on the basis of their belief about the underlying cause of peptic ulcers, and their beliefs about the underlying cause were influenced by the outcome of these different treatments.

## 5 Conclusion

Preventing failed learning in this model is very similar to the problem of maintaining what Kitcher calls “the division of cognitive labor” (1990, 1993). This is the problem of encouraging scientists to work on theories they believe to be duds in order to secure an optimal community response. Maintaining this division of labor prevents the abandonment of optimal theories when experimental results are misleading or priors are biased. Kitcher’s solution to this problem is to appeal to the economic interests of the scientists by offering rewards to those who pursue other avenues. Kitcher (1990, 1993) and Strevens (2003a, 2003b) suggest that our current method of giving rewards to those who were the first to succeed has this effect.

This solution to the problem has the unfortunate consequence of being incompatible with our theories of good epistemic behavior for individuals. That is, scientists are doing well, under Kitcher’s model, when they are actively pursuing the theory they believe to be incorrect with the hopes of gaining a big reward if the theory turns out to be true. In this paper we have another possible solution to the problem that does not rely on that type of epistemic impurity. Our scientists are genuinely pursuing those projects that they deem to be most likely to succeed, but the division of labor has been maintained sufficiently long by limiting the information available to our scientists.<sup>12</sup>

---

11 Success here can be defined in almost any way you like it. Different methods might have different probabilities of generating true explanations, adequate predictions, useful policy suggestions, etc. My hope here is to develop a model which would appeal to people with varying commitments about what constitutes scientific success.

12 For a more detailed discussion of the cognitive division of labor, and the relation of this model to those problems see (Zollman 2009).

Even beyond the problem of maintaining the division of cognitive labor, this model suggests that in some circumstances there is an unintended benefit from scientists being uninformed about experimental results in their field. This is not universally beneficial, however. In circumstances where speed is very important or where we think that our initial estimates are likely very close to the truth, connected groups of scientists will be more reliable. On the other hand, when we want accuracy above all else, we should prefer communities made up of more isolated individuals.<sup>13</sup>

---

<sup>13</sup> The author would like to thank Brian Skyrms, Kyle Stanford, Jeffrey Barrett, Bruce Glymour, and the participants in the Social Dynamics Seminar at UCI for their helpful comments. Generous financial support was provided by the School of Social Science and Institute for Mathematical Behavioral Sciences at UCI.

## References

- Bala, Venkatesh and Sanjeev Goyal (1998), "Learning from neighbours", *Review of Economic Studies* 65, 565–621.
- Berry, D. A. and B. Fristedt (1985). *Bandit Problems: Sequential Allocation of Experiments*. London: Chapman and Hall.
- Bovens, Luc and Stephan Hartmann (2003), *Bayesian Epistemology*. Oxford: Oxford University Press.
- Ellison, Gregory and Drew Fudenberg (1995), "Word-of-mouth communication and social learning", *The Quarterly Journal of Economics* 110(1), 93–125.
- Goldman, Alvin (1999), *Knowledge in a Social World*. Oxford: Clarendon Press.
- Goldman, Alvin (2009), "Systems-Oriented Social Epistemology", in T. Gendler and J. Hawthorn (eds), *Oxford Studies in Epistemology* vol. 5. New York: Oxford University Press. Reprinted in THIS VOLUME under the title "A Guide to Social Epistemology".
- Herron, Timothy, Teddy Seidenfeld, and Larry Wasserman (1997), "Divisive Conditioning: Further Results on Dilation", *Philosophy of Science* 64, 411–444.
- Hull, David (1988), *Science as a Process*. Chicago: University of Chicago Press.
- Kitcher, Philip (1990), "The division of cognitive labor." *The Journal of Philosophy* 87(1), 5–22.
- Kitcher, Philip (1993), *The Advancement of Science*. New York: Oxford University Press.
- Laudan, Larry (1996), *Beyond Positivism and Relativism: Theory, Method, and Evidence*. Boulder: Westview Press.
- Popper, Karl (1975), "The Rationality of Scientific Revolutions" in *Problems of Scientific Revolution: Progress and Obstacles to Progress*, Oxford: Clarendon Press.
- Rousseau, J. J. (1791 [1954]). *The Social Contract*. New York: Hafner Publishing Company.
- Strevens, Michael (2003a), "Further properties of the priority rule", *Manuscript*
- Strevens, Michael (2003b), "The role of the priority rule in science," *Journal of Philosophy* 100(2), 55–79.
- Zollman, Kevin J.S. (2009) "The Epistemic Benefit of Transient Diversity" *Erkenntnis*, forthcoming.