

The Communication Structure of Epistemic Communities

Kevin J. S. Zollman^{†‡}

Increasingly, epistemologists are becoming interested in social structures and their effect on epistemic enterprises, but little attention has been paid to the proper distribution of experimental results among scientists. This paper will analyze a model first suggested by two economists, which nicely captures one type of learning situation faced by scientists. The results of a computer simulation study of this model provide two interesting conclusions. First, in some contexts, a community of scientists is, as a whole, more reliable when its members are less aware of their colleagues' experimental results. Second, there is a robust tradeoff between the reliability of a community and the speed with which it reaches a correct conclusion.

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 (cf. Goldman 1999; Bovens and Hartmann 2003).

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? For example, Kitcher (1990, 1993) and Strevens (2003a, 2003b) have recently looked at the impact that dif-

[†]To contact the author, please write to: Department of Philosophy, Baker Hall 135, Carnegie Mellon University, Pittsburgh, PA 15213-3890; e-mail: zkollman@andrew.cmu.edu.

[‡]The author would like to thank Brian Skyrms, Kyle Stanford, Jeffrey Barrett, Bruce Glymour, and the participants in the Social Dynamics Seminar at University of California–Irvine for their helpful comments. Generous financial support was provided by the School of Social Science and Institute for Mathematical Behavioral Sciences at UCI.

Philosophy of Science, 74 (December 2007) pp. 574–587. 0031-8248/2007/7405-0003\$10.00
Copyright 2007 by the Philosophy of Science Association. All rights reserved.

ferent 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.

These two projects need not compete with one another. While it is possible that the best epistemic communities are made up of epistemically “sullied” individuals, we have no *a priori* reason to think this is the case.¹ Neither is it the case that a theory of proper individual epistemic conduct answers all the question of community design. Once one fully articulates a theory of individual epistemic rationality, it is still an open question what the optimal community structure is for these agents—the individualistic question is only part of the answer.

A community is made up of many facets, and there are many questions to be answered when it comes to optimal epistemic communities. Here we will be interested in one feature of communities, 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, Bala and Goyal (1998). The surprising result of this analysis is that in many cases a community made up of less informed individuals is more reliable at learning correct answers. Reducing information to scientists, one might expect 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 tradeoff 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. Following that the limitations of the model as a model of science are discussed in Section 4, 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 method of treatment which 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, work on it will not result in any new in-

1. Some have suggested that the goals of epistemically virtuous communities may conflict with the goal of epistemically virtuous individuals, primarily in discussing a community virtue recently called the “division of cognitive labor” (cf. Popper 1975; Hull 1988; Kitcher 1990, 1993). While this certainly may be the case, more work is needed to demonstrate that this virtue cannot be achieved in other ways.

formation 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*, assign the following probabilities to the superiority of the new treatment: 0.33, 0.49, 0.51, and 0.66. They then each pursue the treatment method which they think best. Two scientists, *C* and *D*, will pursue the new treatment option and two, *A* and *B*, the old. 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. Specifically suppose all agree on these probabilities:

$$P(\text{The result of } C\text{'s experiment} \mid \text{New method is better}) = 0.4$$

$$P(\text{The result of } D\text{'s experiment} \mid \text{New method is better}) = 0.4$$

$$P(\text{The result of } C\text{'s experiment} \mid \text{New method is worse}) = 0.6$$

$$P(\text{The result of } D\text{'s experiment} \mid \text{New method is worse}) = 0.6$$

After meeting and reporting their results to each other *A*, *B*, *C*, and *D* now assess the probability of the new theory being better as 0.1796, 0.2992, 0.3163, and 0.4632 respectively. As a result, none of them will pursue the new treatment; we have lost a more beneficial treatment forever. 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.

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.²

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.³ 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 availability of exper-

2. 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.

3. 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.

imental 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 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 ϕ_1 and ϕ_2 and two actions A_1 and A_2 . Action A_1 has the same expected return in both states while A_2 's is lower in ϕ_1 and higher in ϕ_2 . 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 take the action which has the highest expected utility given their beliefs. They receive a payoff from their actions which is independently drawn for each player from a common distribution with the appropriate mean. Each agent observes the outcome of his actions and the outcome of *some* others, and then updates his beliefs based on simple Bayesian reasoning about the state of the world.⁴

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. ϕ_1 and ϕ_2 respectively represent the state where the current method and the new method is better. Bala and Goyal endeavor to discover under what conditions correct convergence can be guaranteed. They consider two different restrictions, restrictions on priors and restrictions on information.

The second suggestion, limiting information, will be our primary focus here. This restriction is achieved by limiting which other agents an 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 only to see those agents with which she is directly connected.

Bala and Goyal consider agents arranged on a line where each agent can only see 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, the probability of converging to the wrong result is no longer zero! This is a

4. '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).

remarkable result, because it contradicts a basic intuition about science: that access to more data is always better.⁵ 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 ϕ_1 and ϕ_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 converge. One might wonder, in these cases how much influence the 'royal family' would have on the population. 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. 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 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 is connected to everyone else, Bala and Goyal's royal family. The last network is one where everyone is connected to everyone.

We will, unbeknownst to our agents, set the world in ϕ_2 , where the new

5. Ellison and Fudenberg (1995) present a different model which comes to the same conclusions. In their model, the interaction structure is not fixed, individual 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 for even 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.




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

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 which they are connected. A population of agents is considered 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 which 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 ϕ_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.

The results of a computer simulation are presented in Figures 2 and 3. In Figure 2, the x-axis represents total number of agents and y-axis represents the proportion of 10,000 runs that reached the correct beliefs.⁶ 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.

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 which are more connected. This also sug-

6. Although it is possible for a population to continue unfinished indefinitely, no population failed to converge.




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

gests 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.⁷

So far, we have only looked at the properties of three networks, the trend seems to be that increased connectivity corresponds to faster but less reliable convergence. This is generalizing from three, relatively extreme

7. 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.




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

networks, however. It would be good to engage in a more systematic survey.

3.2. Connectivity and Success. For relatively small sizes (less than seven), we can exhaustively search the properties of all networks. The suggestion in the previous section, that decreased connectivity results in slower, but more reliable learning, can be tested more extensively. In the previous section, connectivity was left as an intuitive criterion. In fact, there are several graph statistics that correspond to our notion of connectivity. Here, we will use *density* which represents the percentage of possible connections that actually obtain in a graph.

Taking all networks (up to isomorphism) between size three and six we can compare these statistics to network's learning properties. These results are presented in Figure 4.⁸ A regression among the largest group (networks of size six) reveals that density is a stronger predictor of successful learning than any other common graph statistic. Only one other graph statistic significantly improves the prediction beyond density alone, that is the

8. This is a result of running 10,000 trials for every graph, up to isomorphism, of size six or lower. The initial beliefs were independently drawn from a uniform distribution over [0, 1].




Figure 4. Density versus probability of successful learning.

clustering coefficient.⁹ This statistic measures the degree to which one's neighbors (those to whom an individual is connected) are connected to each other. In both cases, the lower the statistic (i.e., the less dense and less clustered a graph is) the higher the successful learning. In addition, the in-network degree variance is not correlated with success, suggesting that it is not the centrality of the wheel, but its high connectivity that results in its decrease in reliability.¹⁰

Examining the differences among the different finite cases is instructive. It appears that sparsely connected networks have a much higher 'inertia'. This inertia takes two forms. First, an unconnected network experiences less widespread change 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 less than expectation is four times higher in

9. In fact, there are two clustering coefficients in the literature. One from Newman et al. (2002) is a slightly better predictor than one presented in Watts (1999) although the difference is very small.

10. The in-network degree variance measures the variability of the number of neighbors for each agent in a network. This number will be high when certain agents are connected to more people than the average, thus representing unequal connectivity. Since this statistic is uncorrelated with success, we can conclude that the centrality of the royal family does not influence its reliability.




Figure 5. The five most accurate (top) and five fastest (bottom) networks.

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 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.¹¹

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.

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 5). Four of the five most reliable graphs are minimally connected—that is, 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

11. The statistics reported here are comparing 100 runs of complete six-person networks, with the most reliable six-person network pictured in Figure 5.




Figure 6. Speed versus accuracy for networks of size six.

from complete graphs. Figure 6 compares the average time to success and probability of success for 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.

This confirms the tradeoff suggested before, in order to gain the reliability that limiting information provides, one must sacrifice other benefits, in this case, speed. In fact, the tradeoff 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.

In the previous section, relationship between speed and size was a strange one. For complete networks, as the network grew the average time to success of these groups decreased. On the other hand, for wheels and cycles as the network grew the average time to success increased. This diversity is verified by the more complete analysis. There appears to be no correlation between size and speed when all networks are considered.

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 tradeoffs do exist and that one can achieve increased reliability by limiting information.

4. The Right Model. There are three assumptions that underlie this model which might cause some concern. They are:

1. The learning in our model is governed by the observation of payoffs.
2. There is a 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.

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.¹²

The second and third assumptions are less innocuous. Similar conclusion can be reached by analyzing another model which results from relaxing these assumptions. Unfortunately, space prohibits a discussion of these results here. It should not be presumed, however, that the Bala-Goyal model is inapplicable. In fact, this model very closely mimics 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 learning circumstances

12. This is not to say that true theories always have higher payoffs. Instead, this model is so general as to apply to either circumstance.

that conform to these three assumptions. 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). Experimental techniques and apparatus may well follow a similar pattern.

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. Here we have another possible solution to the problem which does not rely on that type of epistemic impurity. Our scientist are genuinely pursuing those projects which 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.

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 scientist 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.

REFERENCES

- Bala, Venkatesh, and Sanjeev Goyal (1998), “Learning from Neighbours”, *Review of Economic Studies* 65: 565–621.
- 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”, *Quarterly Journal of Economics* 110: 93–125.
- Goldman, Alvin (1999), *Knowledge in a Social World*. Oxford: Clarendon.

- 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", *Journal of Philosophy* 87: 5–22.
- (1993), *The Advancement of Science*. New York: Oxford University Press.
- Laudan, Larry (1996), *Beyond Positivism and Relativism: Theory, Method, and Evidence*. Boulder, CO: Westview.
- Newman, M. E. J., D. J. Watts, and H. J. Strogatz (2002), "Random Graph Models of Social Networks", *Proceedings of the National Academy of Sciences* 99: 2566–2572.
- Popper, Karl (1975), "The Rationality of Scientific Revolutions", in Rom Harré (ed.), *Problems of Scientific Revolution: Progress and Obstacles to Progress*. Oxford: Clarendon, 72–101.
- Strevens, Michael (2003a), "Further Properties of the Priority Rule", manuscript. <http://www.strevens.org/research/scistruc/MorePrior.pdf>.
- (2003b), "The Role of the Priority Rule in Science", *Journal of Philosophy* 100: 55–79.
- Watts, Duncan (1999), *Small Worlds: The Dynamics of Networks between Order and Randomness*. Princeton, NJ: Princeton University Press.