

The Epistemic Benefit of Transient Diversity

Kevin J. S. Zollman

Received: 7 October 2008 / Accepted: 28 September 2009 / Published online: 22 October 2009
© Springer Science+Business Media B.V. 2009

Abstract There is growing interest in understanding and eliciting division of labor within groups of scientists. This paper illustrates the need for this division of labor through a historical example, and a formal model is presented to better analyze situations of this type. Analysis of this model reveals that a division of labor can be maintained in two different ways: by limiting information or by endowing the scientists with extreme beliefs. If both features are present however, cognitive diversity is maintained indefinitely, and as a result agents fail to converge to the truth. Beyond the mechanisms for creating diversity suggested here, this shows that the real epistemic goal is not diversity but transient diversity.

A striking social feature of science is the extensive division of labor. Not only are different scientists pursuing different problems, but even those working on the same problem will pursue different solutions to that problem. This diversity is to be applauded because, in many circumstances one can simply not determine a priori if a general theoretical or methodological approach will succeed without first attempting to apply it, study the effects of its application, and develop additional auxiliary theories to assist in its application.¹

The value of diversity presents a problem for more traditional approaches to scientific methodology, since if everyone employs the same standards for induction and has access to the same information, we ought to expect them to all adopt the

¹ When defending alternatives to the classic interpretation of quantum mechanics, Feyerabend succinctly declares, “It takes time to build a good theory” (1968, p. 150).

theory which at the current time looks most promising. Kuhn suggests this problem represents a failure of the traditional approach,

Before the group accepts [a scientific theory], a new theory has been tested over time by research of a number of [people], some working within it, others within its more traditional rival. Such a mode of development, however, *requires* a decision process which permits rational men to disagree, and such disagreement would be barred by the shared algorithm which philosophers have generally sought. If it were at hand, all conforming scientists would make the same decision at the same time (Kuhn 1977, p. 332).

Kuhn's approach to the problem, allowing diversity in standards for induction, is supported by others who offer similar solutions (Hull 1988; Sarkar 1983; Solomon 1992, 2001).

Alternatively, Philip Kitcher (1990, 1993, 2002) and Michael Strevens (2003a, b) have both suggested that homogeneity in inferential strategy can still produce diversity if the scientific reward system has an appropriate structure and scientists are appropriately motivated. In a similar vein, Paul Thagard (1993) has suggested that a uniform method but differential access to information can be of some assistance in maintaining this diversity.

Just as traditional epistemologists ignored the benefit of diversity, many of these contemporary champions of diversity ignore the method by which ultimate consensus is achieved.² They do however point to an important learning situation faced by scientists. In these situations, information about the effectiveness of a theory or method can only be gathered by scientists actively pursuing it. But scientists also have some interest in pursuing a theory which turns out to be right, since effort developing an inferior theory is often regarded as a waste. This is precisely the circumstance described by Kuhn.

Rather than focusing on diversity directly, we will consider this type of learning situation as a problem in social epistemology. We will begin this investigation by presenting an important episode from the history of science which illustrates the need for diversity. After presenting the history of the investigation of peptic ulcer disease in Sect. 1, we will investigate a model which captures some of the central features of this episode in the history of science.

This model represents one type of learning situation discussed by the contemporary champions of diversity, and so presents a situation where the benefit of diversity can be explicitly analyzed. The analysis of this model demonstrates a consistent theme, that a certain amount of diversity provides some benefit to the community. One way to attain this diversity is by limiting the amount of information available to the scientists. This is achieved by arranging them so that they only see a proper subset of the total experiments performed, using different

² Feyerabend is perhaps the most extreme in this regard. He says, "This plurality of theories must not be regarded as a preliminary stage of knowledge that will at some time in the future be replaced by the 'one true theory.'" (1965, p. 149). But even in papers of this era (1965, 1968), he does not advocate holding onto inferior theories indefinitely.

subsets for each agent. This preserves the diversity present in the agents' priors which helps facilitate exploration.

Diversity can be maintained in the models in a second way, extreme priors. If agents have very extreme beliefs, this has the same effect as limiting information. But, this solution offers an opportunity to see the downside to diversity. If agents are both extreme in their initial beliefs and limited in their access to information, the initial diversity is never abandoned. This shows that in these learning situation diversity is not an independent virtue which ought to be maintained at all costs, but instead a derivative one that is only beneficial for a short time.³

Beyond simply maintaining diversity, we also find that in some cases apparently irrational individual behavior may paradoxically make communities of individuals more reliable. In this case we have shown that the socially optimal structure is one where scientists have access to *less* information. But, before looking at the model in detail, we will first turn to an actual case of harmful homogeneity in science.

1 The Case of Peptic Ulcer Disease

In 2005, Robin Warren and Barry Marshall received the Nobel Prize in Physiology or Medicine for their discovery that peptic ulcer disease (PUD) was primarily caused by a bacteria, *Helicobacter pylori* (*H. pylori*). The hypothesis that peptic ulcers are caused by bacteria did not originate with Warren and Marshall, it predates their births by more than 60 years. But, unlike other famous cases of anticipation, this theory was the subject of significant scientific scrutiny during that time. To those who have faith in the scientific enterprise, it should come as a surprise that the widespread acceptance of a now well supported theory should take so long. If the hypothesis was available and subjected to scientific tests, why was it not widely accepted long before Warren and Marshall? Interestingly, the explanation of this error does not rest with misconduct or pathological science, but rather with simple, and perhaps unavoidable, good faith mistakes which were convincing to a wide array of scientists.

The bacterial hypothesis first appeared in 1875. Two bacteriologists, Bottcher and Letulle argued that peptic ulcers were caused by an unobserved bacteria. Their claim was supported by observations of bacteria-like organisms in glands in the stomach by another German pathologist. Almost simultaneously, the first suggestions that PUD may be caused by excess acid began to appear (Kidd and Modlin 1998).

Although both hypotheses were live options, the bacterial hypothesis had early evidential support. Before the turn of the century, there were at least four different observations of spirochete organisms (probably members of the *Helicobacter* genus) in stomachs of humans and other mammals. Klebs found bacteria in the gastric glands in 1881 (Fukuda et al. 2002), Jaworski observed these bacteria in sediment

³ This result is explicitly impossible in the models considered by Kitcher and Strevens. In their models a theory succeeds or fails and this success or failure is known by all agents.

washings in 1889 (Kidd and Modlin 1998), Bizzozero observed spiral organisms in dogs in 1892 (Figura and Bianciardi 2002), and Saloon found similar spirochetes in the stomachs of cats and mice in 1896 (Buckley and O'Morain 1998).

By the turn of the century, experimental results appeared to confirm the hypothesis that a bacterial infection might be, if not an occasional cause, at least an accessory requirement for the development of gastroduodenal ulcers. Thus, although a pathological role for bacteria in the stomach appeared to have been established, the precise role of the spirochete organisms remained to be further evaluated. (Kidd and Modlin 1998)

During the first half of the twentieth century, it appeared that both hypotheses were alive and well. Observations of bacteria in the stomach continued, and reports of successful treatment of PUD with antibiotics surfaced.⁴ At the same time the chemical processes of the stomach became better understood, and these discoveries began to provide some evidence that acid secretion might play a role in the etiology of PUD. Supporting the hypoacidity theory, antacids were first used to successfully reduce the symptoms PUD in 1915. This success encouraged further research into chemical causes (Buckley and O'Morain 1998). In 1954, a prominent gastroenterologist, Palmer, published a study that appeared to demonstrate that no bacteria is capable of colonizing the human stomach. Palmer looked at biopsies from over 1,000 patients and observed no colonizing bacteria. As a result, he concluded that all previous observations of bacteria were a result of contamination (Palmer 1954).

The result of this study was the widespread abandonment of the bacterial hypothesis, poetically described by Fukuda, et al.,

[Palmer's] words ensured that the development of bacteriology in gastroenterology would be closed to the world as if frozen in ice... [They] established the dogma that bacteria could not live in the human stomach, and as a result, investigation of gastric bacteria attracted little attention for the next 20 years (Fukuda et al. 2002, pp. 17–20).

Despite this study, a few scientists and clinicians continued work on the bacterial hypothesis. John Lykoudis, a Greek doctor, began treating patients with antibiotics in 1958. By all reports he was very successful. Despite this, he was unable to either publish his results or convince the Greek authorities to accept his treatment. Undeterred, he continued using antibiotics, an action for which he was eventually fined (Rigas and Papavassiliou 2002). Although other reports of successful treatment with antibiotics or observation of bacteria in the stomach occasionally surfaced, the bacterial hypothesis was not seriously investigated until the late 1970s.

At the 1978 meeting of the American Gastroenterology Association, it appeared that tide had begun to turn. At this meeting, it was widely reported that the current acid control techniques could not cure ulcers but merely control them (Peterson et al. 2002). When antacid treatment was ceased, the symptoms would invariably return. The very next year, Robin Warren first observed *Helicobacters* in a human

⁴ Although, bismuth (an antimicrobial) had been used to treat ulcers dating as far back as 1868, the first report of an antibiotic occurs in 1951 (Unge 2002).

stomach, although reports of this result would not appear in print until 1984 (Warren and Marshall 1984).

Initial reactions to Warren and Marshall's discovery were negative, primarily because of the widespread acceptance of Palmer's conclusions. Marshall became so frustrated with his failed attempts to convince the scientific community of the relationship between *H. pylori* and PUD, that he drank a solution containing *H. pylori*. Immediately after, he became ill and was able to cure himself with antibiotics (Marshall 2002). Eventually after replication of Marshall and Warren's studies, the scientific community became convinced of Palmer's error. It is now widely believed that *H. pylori* causes PUD, and that the proper treatment for PUD involves antibiotics.

While we may never really know if Palmer engaged in intentional misconduct, the facts clearly suggest that he did not.⁵ Palmer failed to use a silver stain when investigating his biopsies, instead relying on a Gram stain. Unfortunately, *H. pylori* are most evident with silver stains and are Gram negative, meaning they are not easily seen by using the Gram stain. Although the silver staining technique existed in the 1950s, it would have been an odd choice for Palmer. That stain was primarily used for neurological tissue and other organisms that should not be present in the stomach. Warren did use the silver stain, although it is not clear what led him to that choice.

Without looking into the souls of each and every scientist working on PUD, one can hardly criticize their behavior. They became aware of a convincing study, carefully done, that did not find bacteria in the stomach. Occasionally, less comprehensive reports surfaced, like those of Lykoudis, but since they contradicted what seemed to be much stronger evidence to the contrary, they were dismissed. Had the acid theory turned out to be true, the behavior of each individual scientist would have been laudable.

Despite the fact that everything was "done by the book," so to speak, one cannot resist the urge to think that perhaps things could have been done differently. In hindsight, Palmer's study was too influential. Had it not been as widely read or been as convincing to so many people, perhaps the bacterial theory would have won out sooner. It was the widespread acceptance of Palmer's result which led to the premature abandonment of the diversity in scientific effort present a few years earlier. But it is not as if we would prefer scientists to remain diversified forever, today we consider further effort attempting to refute the bacterial hypothesis wasteful. Thinking just about PUD, these are likely idle speculations. On the other hand we might consider this type of problem more generally. We might then ask, what features of individual scientists and scientific communities might help to make these communities less susceptible to errors like Palmer's.⁶

⁵ Marshall speculates that the long delay between the reports of his own discovery and the widespread acceptance of the bacterial hypothesis were (partially) the result of the financial interests of pharmaceutical companies (Marshall 2002). While this may be an example of pathological science, the dismissal of the bacterial hypothesis from 1954 to 1985 probably is not.

⁶ For the interested reader, Paul Thagard (1998a, b) presents a discussion of the relationship between this episode and the history of science which considers different philosophical aspects of the history.

In order to pursue this more general analysis, we will turn to abstract models of scientific behavior. Through the analysis of these models, perhaps we can gain insight into ways that scientific communities may be more and less resistant to errors of this sort.

2 Modeling Social Interactions

In attempting to generate a more abstract model, we must both find a way to represent the learning problem faced by PUD scientists and also a way to model the communication of their results. We will here generalize an approach described in (Bala and Goyal 1998; Zollman 2007). This model employs a type of problem devised in statistics and used in economics known as bandit problems and models communication among scientists as a social network (a technique widely used in the social sciences).

2.1 Bandit Problems: The Science of Slot Machines

Beginning in the 1950's there was increasing interest in designing statistical methods to more humanely deal with medical trials, like those conducted for PUD (see, e.g. Robbins 1952). A scientist engaging in medical research is often pulled in two different directions by her differing commitments. On the one hand she would like to gain as much information as possible, and so would prefer to have two large groups, one control and one experimental group. On the other hand, she would like to treat as many patients as possible, and if it appears that her treatment is significantly better than previous treatments she might opt to reduce the size of the control group or even abandon the experiment altogether.⁷ These two different concerns lead to the development of a class of problems now widely known as bandit problems.

The underlying metaphor is that of a gambler is confronted with two slot machines which payoff at different rates. The gambler sequentially chooses a slot machine to play and observes the outcome. For simplicity we will assume that the slot machines have only two outcomes, "win" or "lose".⁸ The gambler has no idea about the probability of securing a "win" from each machine; he must learn by playing. Initially it seems obvious that he ought to try out both machines in order to secure information about the payoffs of the various machines, but when should he stop? The controlled randomized trial model would have him pull both machines equally often in order to determine with the highest accuracy the payoffs of the two machines. While this strategy would result in the most reliable estimates at its conclusion, it will often not be the most remunerative strategy for the gambler.

⁷ For example, a recent study on the effect of circumcision on the transmission of HIV was stopped in order to offer circumcision to the control group because the effect was found to be very significant in the early stages of research.

⁸ Nothing requires that we limit ourselves in this way. The models described later in this paper have a larger set of possible outcomes.

The gambler is confronted with a problem that is identical to the clinical researchers problem, he wants to gain information but he would also like to play the better machine. Discovering the optimal strategy in these circumstances is very difficult, and there is significant literature on the subject (see, e.g. Berry and Fristedt 1985).

While used primarily as a model for individual clinical trials, situations of this structure pervade science. PUD researchers were confronted with this choice when they decided which of two different treatment avenues to pursue. They could dedicate their time to developing more sophisticated acid reduction techniques or they could search for the bacteria that might cause PUD. Their payoff, here, is the reward given for developing a successful treatment, and their probability of securing that reward is determined largely by whether or not a bacteria is responsible for PUD. Since Palmer's results suggested that this was unlikely, the researchers believed they were more likely to secure success by pursuing research relating to acid reduction.

The application of other theories may follow a similar pattern. In population biology for instance, one can pursue many different modeling techniques. One might choose to develop in more detail a model of a particular sort, say a finite population model, because one believes that this type of model is more likely to yield useful results. Of course, one cannot always know a priori whether a type of model will be successful until one attempts to pursue it. Again, we have a circumstance akin to a bandit problems. Zollman (2007) suggests that the bandit problem model fits well with Laudan's (1996) model of paradigm change, since Laudan believes these changes are based on something like expected utility calculations.

Here the payoff of the bandit is analogous to a successful application of a given theory. One attempts to apply a theory to a particular case and succeeds to different degrees. Individual theories have objective probabilities of success which govern the likelihood that it can be successfully applied in some specified domain. Since they are interested in successful applications of a theory, scientist would like to work only on those theories which can most successfully be applied. It is usually assumed in the bandit problem literature that the payoff to a particular bandit (or treatment, or theory, etc.) is an independent draw from a distribution and this distribution remains constant over time. That is, past success and failure do not influence the probability of success on this trial (conditioning, of course, on the underlying distribution).

Are scientific theories like this? Do past successes influence the probability of future success? Certainly they do, and they can in many different ways. For instance, it may be that a particular theory has only finitely many potential successful applications. Applying the theory then is like drawing balls from an urn without replacement. Disproportionate past success now reduces the chance of future success, because most of the successful applications have been found.⁹ It might also be the case that past success increases the chance of future success, a new successful application of a scientific theory might open up many new avenues for

⁹ This point is due to a conversation with Michael Weisberg and Ryan Muldoon. They present a rather different model of scientific practice which centrally models science in this way (Weisberg and Muldoon 2008).

application that were not previously available. In this case, previous success increases the probability of future success.

There are many cases however, where the assumption of independence is not far off. Even if there are only finitely many possible applications, if that number is very large the probability of current success will be very near to the probability of past success. Alternatively, perhaps the number of new potential applications opened up by a past success balances out the number of past successes. While I acknowledge the limitations of the assumption, this model will assume that individual attempts at applications are independent and drawn from a common distribution for that scientific theory. Extending this model to cases where the probability of future success is determined by past success will be left to future research.

In this model we will present the scientists with the choice between two potential methods (two “bandits”). A scientist will choose a method an attempt to apply it (a pull of a bandit’s arm) Success will be represented by a draw from a binomial distribution ($n = 1,000$). This number represents the degree of success of that method (the payoff from a pull of a bandit’s arm). Each method has a different intrinsic probability of success, and scientists pursue the theory they think is currently most likely to succeed on each given application. They update their beliefs using Bayesian reasoning (described below) based on their own success and also based on the success of some others. In order to allow us to study different ways of distributing the results, we will use social networks to represent which experimental results are observed by each individual.

The reader will note that this model of scientific practice is a bit different from the traditional ones. Loosely following in the pragmatist tradition, scientists are not passive observers of evidence, instead they are actively engaged in the process of evidence gathering.¹⁰ In addition, rather than gathering evidence for or against a particular theory, they are attempting to estimate the efficacy of different methodologies which can both succeed but to differing degrees. I do not mean to suggest that the traditional model is wrong—scientists are often engaged in both types of inquiry. Instead, I present this model as one of many which treats a particular aspect of scientific practice.

2.2 Social Networks

Beginning in sociology and social psychology, scholars have become increasingly interested in mathematically representing the social relations between people. Much of recent work on social networks in sociology has focused on studying the underlying structure of various social networks and modeling the evolution of such networks over time. However, there is also increasing interest in considering how existent social networks effect the change of behaviors of individuals in those networks. Two recent examples include Alexander’s (2007) work on the evolution of strategies in games and Bala and Goyal’s (1998, see also Goyal 2005) work on

¹⁰ Although evidence is always arriving, what evidence arrives depends on the actions taken by individual scientist. Their actions, depend, in turn, on what their beliefs are about the efficacy of different methodologies. This represents a violation of the “relative autonomy of facts” which is criticized by Feyerabend (1965, 1968).

learning in networks. Both of these studies have found that different network structures can have significant influence on the behavior of the system.

A social network consists of a set of individuals (nodes) connected to one another by edges. Edges may be directed or undirected depending on whether or not the underlying relationship being represented is symmetric. While edges in a social network can represent many different relationships, we will here focus on the transmission of information from one person to another. A link between two individuals represents the communication of results from each one to the other. We will here presume that this relationship is symmetric, and so will use undirected graphs.

2.3 Learning in Bandit Problems

As suggested above, individuals will learn based on Bayesian reasoning. In earlier work, I (Zollman 2007) studied the effect of social interaction in Bandit problems with a very limited set of potential outcomes. I considered a circumstance where the probability of success of one methodology was known—learning outcomes from it were uninformative—while the other was either x or y (with x being higher than the mean outcome of the known action and y being lower). This model may not be sufficiently general; scientist can entertain a large variety of possibilities which are informed by previous success and failure.

If they entertain a large enough range of alternatives we cannot use the simple Bayesian model of discrete hypotheses. Instead we must turn to using functions to represent an agent’s belief over infinitely many hypotheses. One way of modeling this type of Bayesian learning for Bernoulli trials is with beta distributions.

Definition 1 (*Beta Distribution*) A function on $[0, 1], f(\cdot)$, is a beta distribution iff for some $\alpha > 0$ and $\beta > 0$

$$f(x) = \frac{x^{\alpha-1}(1-x)^{\beta-1}}{B(\alpha, \beta)}$$

where $B(\alpha, \beta) = \int_0^1 u^{\alpha-1}(1-u)^{\beta-1} du$.

The beta distribution has some very nice properties. In addition to having only two free parameters it has the property that if one has a beta distribution as a prior, and one takes a sample of any size and updates, one will have a beta distribution as a posterior. Suppose someone has a coin of unknown bias and performs n flips of the coin and receives s heads. If this individual has a beta distribution with parameters α and β as a prior, then his posterior will also be a beta distribution with the posterior parameters $\alpha + s$ and $\beta + n - s$ (cf. DeGroot 1970).

The expectation for a beta distribution is given by $\frac{\alpha}{\alpha+\beta}$. This enables a rather brief convergence result. Suppose an agent starts out with priors α and β . She then performs a series of trials which result in s successes in n trials. Her posterior has parameters $\alpha + s$ and $\beta + n - s$. As a result her posterior mean is given by:

$$\frac{\alpha + s}{\alpha + s + \beta + n - s} = \frac{\alpha + s}{\alpha + \beta + n}$$




Fig. 1 Learning using beta distributions

As s and n grow, $\frac{s}{n}$ approaches the true probability of successes and since they will eventually grow well beyond α and β , the mean of the agent's beliefs will approach the true mean (cf. Howson and Urbach 1996).

In addition, learning via beta distributions is relatively efficient. Figure 1 illustrates simulation results for a single agent learning the probability of a coin. Her starting distribution is a beta distribution where α and β are randomly chosen to be any number between zero and four. She then flips the coin 10 times and then updates her priors.¹¹ The x-axis represents the number of such experiments that have been performed. The lines represent the average distance from the true mean for 100 such individuals. Here we see that beta distribution learning can be very fast.

2.4 Individual Choice

Bandit problems represent a difficult problem for the player. If one knows that one has many pulls available and the present is no more important than the distant future, one should be willing to (at least sometimes) pull the arm that one regards as inferior in order to ensure that one is not mistaken about which arm is best. However, if one has only a few pulls or if today's pull is much more important than tomorrow's, then one may not be willing to explore when it comes at some cost. Significant mathematics have been developed in order to determine what sorts of trade-offs one should be willing to make. Determining one's optimal strategy can become complex even in relatively simple circumstances (Berry and Fristedt 1985).

The movement from a single gambler to a group of gamblers who can observe one another introduces yet another layer of complexity. If you and I are both playing slot machines and you get to observe how I do, you may want to play the machine you currently think is better and leave the exploration to me. After all, you learn as

¹¹ The chose of $[0, 4]$ was chosen so that the initial beliefs do not swamp even a single experimental result.

much from the information I get as the information you get, but only you lose by playing the machine that seems currently worse. This introduces what is known as a free-rider problem in economics. We both want to leave the exploration to the other person, and so it is possible that no one will do it.

For various reasons, we will ignore these complexities, and instead focus on individual scientists who are *myopic*. They are unwilling to play a machine they think inferior in order to gain the information from that machine. Instead, on each round they will pull the arm that the currently think is the best arm, without regard for the informational value of pulling the inferior one. Why should we make this assumption? First, I think it more closely accords with how individual scientists choose methodologies to pursue. Second, it closely mimics what would happen if individuals were being selfish and leaving the exploration to one another. This is a possibility in social circumstances because of the free-rider problem noted above. Third, myopic behavior is optimal when one cares significantly more about the current payoff than about future payoffs. Scientists are rewarded for current successes (whether it be via tenure, promotion, grants, or awards), and I believe that this causes a sort of myopia *with large scale decisions* like research methodology. Finally, even if one finds this assumption unreasonable, it represents an interesting starting point from which we can gauge the effect of making scientists care increasingly about the future. This is another avenue of research that has yet to be explored, but could provide interesting insights.

3 Limiting Information

Individuals will be assigned a random initial α_i 's and β_i 's for each action i from the interval $[0, 4]$. Since there is more than one available methodology which might be applied, it is possible for individuals to lock-in on the sub-optimal methodology. Consider a single learner case where the individual has the following priors:

$$\begin{aligned}\alpha_1 &= 1 \\ \beta_1 &= 3 \\ \alpha_2 &= 3 \\ \beta_2 &= 1\end{aligned}$$

This yields an expectation of 0.25 for action 1 and 0.75 for action 2. Suppose that action 2 does have an expectation of 0.75, but action 1 has a higher objective expectation of 0.8. Since the individual thinks that action 2 is the superior action, he will take it on the first round. So long as the results he receives do not take him too far from the expectation of action 2, he will continue to believe that action 2 is superior and will never learn that his priors regarding action 1 are skewed.

Of course estimating the probability of this happening can be very complex, especially in social settings where individuals' priors are influenced by other individuals who are also influenced by still other individuals. Instead of attempting to prove anything about this system, we will simulate its behavior. In these simulations a group of individuals is allowed to successively choose a methodology to apply and learn the degree of success and the method chosen by their neighbors.




Fig. 2 A 10 person cycle, wheel, and complete graph

After 10,000 iterations, we observe whether every individual has succeeded in determining which methodology has the highest intrinsic probability of success.

We will first consider three networks which represent idealizations of different social circumstances (the networks pictured in Fig. 2). The first network, the cycle, represents a circumstance of perfect symmetry—every individual has exactly two neighbors and each is equally influential. In this network, information is not widely shared. The network on the right, the complete graph, represents the other extreme where all information is shared. Again, like the cycle, there is perfect symmetry, but the amount of information is radically different. The middle graph, the wheel, represents a circumstance where symmetry is broken. Here the individual at the center shares his information with everyone, but other individuals do not.

Results for the simulation on these three networks is presented in Fig. 3.¹² Here we find a surprising result, information appears to be harmful. The cycle—where each individual has access to the smallest amount of information—is superior, followed by the wheel and then by the complete graph. The degree of difference here should 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.

It would appear here that the amount of information distributed is negatively impacting the ability of a social group to converge on the correct methodology. Initially suggestive information is causing everyone to adopt one particular methodology. Because of the stochastic nature of method application, sometimes an inferior method can initially appear fruitful or a superior one might seem hopeless. When this happens, in the complete graph, the information is widely disseminated and everyone adopts the inferior theory. Once adopted, they are no longer learning about the other methodology, and unless the right sort of results occur, they will not ever return. This illustrates how central the assumption that evidence be generated by applications of a given methodology (the violation of Feyerabend's "relative autonomy of facts"). If the evidence available to the scientists did not depend on their own choices, this problem could not come about. Here we have a situation much like the one of peptic ulcer disease discussed above. The widely read study of Palmer suggested that one particular method was unlikely to succeed. Because of its

¹² Each trial represents 1,000 "pulls", which have a 0.5 and 0.499 probability of "winning" respectively.




Fig. 3 Results for the cycle, wheel, and complete graph




Fig. 4 Results for all networks of size six

influence, everyone began researching the other method and did not learn about the effectiveness of the former. It was not until repeated failures of the acid suppression method combined with a remarkable success of the bacterial, that individuals were willing to entertain switching methods.

In order to ensure that our generalization from these three networks is not due to some other feature of the networks, we can consider a wider search. It is possible to exhaustively search all networks with up to six individuals and compare their relative probabilities of success. Figure 4 shows the result of this simulation. The x-axis represents the density of the graph—the proportion of possible connections which actually obtain in the graph. We see here clearly that the more dense graphs—those where individuals have access to more information—are on average




Fig. 5 Three beta distributions

worse than less dense ones. Again, here it appears that more information is harmful.¹³

4 Different Priors

In beta distributions the size of the initial α 's and β 's determines the strength of an individual's prior belief. Figure 5 illustrates three different beta distributions with differing parameters. Although all three distributions have the same mean, the ones with higher initial parameters have lower variances. Also, individuals who possess these more extreme distributions as priors will be more resistant to initial evidence. If all three receive 8 successes in 12 trials, the expectations after updating will be 0.6, 0.56, and 0.52. The more extreme the initial beliefs an individual has, the more resistant to change she is.

Since failed learning is a result of misleading initial results infecting the entire population, one might suggest that increasing an individual's resistance to change might help to alleviate this danger. In a more informal setting Popper suggested this possibility. "A limited amount of dogmatism is necessary for progress. Without a serious struggle for survival in which the old theories are tenaciously defended, none of the competing theories can show their mettle" (1975, p. 87).¹⁴ If dogmatism could serve this purpose, it might turn out that the difference between less connected and more connected networks would vanish.

In order to investigate this possibility we will observe three canonical networks (a seven person cycle, wheel, and complete graph) and vary the range of initial beliefs. The results in the previous section were for α and β values between zero and

¹³ These results are qualitatively similar to the results obtained by Zollman (2007) in studying a more limited model. This lends additional support to the conclusion that bandit problem-like situations informally is not uniformly helpful.

¹⁴ A similar sentiment is echoed by David Hull (1988, 32) and Miriam Solomon (1992).




Fig. 6 Success versus maximum α and β

four. Figure 6 shows the results for these three graphs as the maximum possible α and β value is increased. Individuals can still have very low initial parameters, but as the maximum increases this becomes less likely.

The results are quite striking. For α 's and β 's drawn from $[0, 1000]$, the results are similar to the smaller initial parameters used before. However, as the maximum grows, the order of the networks reverses itself. At very extreme initial parameters, the complete network is by far the best of the three networks. This is not simply the result of one network reducing its reliability, but rather one network gains while the other loses.

The cause of this reversal is interesting. Complete networks were worse because they learned too fast. The wealth of information available to the agents sometimes caused them to discard a superior action too quickly. In the more limited networks this information was not available and so they did not jump to conclusions.

Similarly, when our agents have very extreme priors, even rather large amounts of information will not cause them to discard their prior beliefs. As a result, the benefit to disconnected networks vanishes—no matter how connected the network, agents will not discard theories too quickly. This explains why the complete network is not so bad, but does not explain why the less connected network becomes worse. In the less connected network there is simply not enough information to overcome the extremely biased priors within sufficient time. Since the simulations are stopped after a certain number of trials, the extreme priors have biased the agents sufficiently that they cannot overcome this bias in time. The drop in reliability of the less connected networks is the result of a substantial increase in networks that failed to unanimously agree on any theory. As 10,000 experimental iterations are probably already an extreme, failing to agree in this time should be judged as a failure.

All of this is illustrated in Fig. 7. Here the top three networks are a seven person cycle, wheel, and complete network with very extreme priors.¹⁵ The bottom three

¹⁵ In this case α and β are drawn from a uniform distribution on $[0, 7000]$.




Fig. 7 In network variance over time

are the same but with priors drawn from a much smaller distribution. The y-axis represents the mean variance of the actions taken on that round. The higher the variance, the higher the diversity in actions.

The steep drop off of the three networks with limited priors is to be expected. The fact that the cycle and wheel preserve their diversity for much longer than the complete network illustrates the benefits of low connectivity. (In fact, the complete network drops off so quickly that it is almost invisible on the graph.) All of the networks with extreme priors maintain their diversity much longer, however, the complete network begins to drop off as more information accumulates.

An interesting feature discovered by varying the extremity in the priors can be seen in Fig. 6. Around a maximum of 3,000 the three networks almost entirely coincide. Here, we are slightly worse off than at either of the extremes (cycle with unbiased priors or complete with extreme priors). On the other hand, this represents a sort of low risk position, since the network structure is largely irrelevant to the reliability of the model.

Returning again to the case of PUD, if individual scientists had been more steadfast in their commitment to the bacterial hypothesis they might not have been so convinced by Palmer's study. Perhaps they would have done studies of their own, attempting to find evidence for a theory they still believed to be true. Because of their steadfast commitments we would want all of the information to be widely distributed so that eventually we could convince everyone to abandon the inferior theory once its inferiority could be clearly established.

5 Conclusion

This last result illustrates an important point. At the heart of these models is one single virtue, transient diversity. This diversity should be around long enough so that individuals do not discard theories too quickly, but also not stay around so long

as to hinder the convergence to one action. One way of achieving this diversity is to limit the amount (and content) of information provided to individuals. Another way is to make individuals' priors extreme. However, these are not independent virtues. Both together make the diversity too stable, and result in a worse situation than either individually.

For PUD, I have suggested that things might have been better had Palmer's result not been communicated so widely or had people been sufficiently extreme in their beliefs that many remained unconvinced by his study. However, it would have been equally bad had both occurred simultaneously. In the actual history, 30 years were wasted by pursuing a sub-optimal treatment. Had the scientists been both uninformed and dogmatic, we might still be debating the bacterial and hypoacidity hypothesis today.

Like Kitcher and Strevens (and contra Kuhn) these models demonstrate that diversity can be maintained despite uniformity in inductive standards. I have shown that, so long as agents have diverse priors, diversity can be maintained by limiting information or by making individuals extreme in their initial estimates. Unlike Kitcher and Strevens, this study provides a series of solutions to the diversity problem which they do not consider. In addition, this study articulates a negative consequence of diversity, a problem which does not occur in Kitcher and Strevens' models.¹⁶

The offered solutions to this problem all turn on individuals being arranged in ways that make each individual look epistemically sub-optimal. The scientists do not observe all of the available information or have overly extreme priors. Looking at these scientists from the perspective of individualistic epistemology, one might be inclined to criticize the scientists' behavior. However, when viewed as a community, their behavior becomes optimal. This confirms a conjecture of David Hull (1988, pp. 3–4), that the characteristic rational features of science are not properties of individuals but instead properties of scientific groups. Seen here, limiting information or endowing individuals with dogmatic priors has a good effect when the overall behavior of the community is in focus. This suggests that when analyzing particular behaviors of scientists (or any epistemic agents) we ought to think not just about the effect their behavior has on their individual reliability, but on the reliability of the community as a whole.

Acknowledgments The author would like to thank Brian Skyrms, Kyle Stanford, Jeffrey Barrett, Bruce Glymour, Sam Hillier, Samir Grover, Kevin Kelly, Teddy Seidenfeld, Michael Strevens, Michael Weisberg, Ryan Muldoon, several contributors at the Wikipedia reference desk, and the anonymous referees for their assistance. Code for the simulations can be obtained from the author's website: <http://www.andrew.cmu.edu/users/kzollman/>.

References

- Alexander, J. M. (2007). *The structural evolution of morality*. Cambridge: Cambridge University Press.
 Bala, V., & Goyal, S. (1998). Learning from neighbours. *Review of Economic Studies*, 65, 565–621.

¹⁶ For a rather different discussion of the two sides of this debate see (Laudan 1984).

- Berry, D. A., & Fristedt, B. (1985). *Bandit problems: Sequential allocation of experiments*. London: Chapman and Hall.
- Buckley, M. J., & O'Morain, C. A. (1998). Helicobacter biology—discovery. *British Medical Bulletin*, 54(1), 7–16.
- DeGroot, M. H. (1970). *Optimal statistical decisions*. NY: McGraw-Hill.
- Feyerabend, P. (1965). Problems of empiricism. In R. G. Colodny (Ed.), *Beyond the edge of certainty. Essays in contemporary science and philosophy* (pp. 145–260). Englewood Cliffs NJ: Prentice-Hall.
- Feyerabend, P. (1968). How to be a good empiricist: A plea for tolerance in matters epistemological. In P. Nidditch (Ed.), *The philosophy of science. Oxford readings in philosophy* (pp. 12–39). Oxford: Oxford University Press.
- Figura, N., & Bianciardi, L. (2002). Helicobacters were discovered in Italy in 1892: An episode in the scientific life of an eclectic pathologist, Giulio Bizzozero. In B. Marshall (Ed.), *Helicobacter pioneers: Firsthand accounts from the scientists who discovered Helicobacters*, (pp. 1–13). Victoria, Australia: Blackwell Science Asia.
- Fukuda, Y., Shimoyama, T., Shimoyama, T., & Marshall, B. J. (2002). Kasai, Kobayashi and Koch's postulates in the history of *Helicobacter pylori*. In B. Marshall (Ed.), *Helicobacter pioneers*, (pp. 15–24). Oxford: Blackwell publishers.
- Goyal, S. (2005). Learning in networks: A survey. In G. Demange & M. Wooders (Eds.), *Group formation in economics: networks, clubs, and coalitions*. Cambridge: Cambridge University Press.
- Howson, C., & Urbach, P. (1996). *Scientific reasoning: The Bayesian approach* (2nd ed). Chicago: Open Court.
- Hull, D. (1988). *Science as a process*. Chicago: University of Chicago Press.
- Kidd, M., & Modlin, I. M. (1998). A century of *Helicobacter pylori*. *Digestion*, 59, 1–15.
- Kitcher, P. (1990). The division of cognitive labor. *The Journal of Philosophy*, 87(1), 5–22.
- Kitcher, P. (1993). *The advancement of science*. New York: Oxford University Press.
- Kitcher, P. (2002). Social psychology and the theory of science. In S. Stich & M. Siegal (Eds.), *The cognitive basis of science*. Cambridge: Cambridge University Press.
- Kuhn, T. S. (1977). Collective belief and scientific change. In *The essential tension* (pp. 320–339). University of Chicago Press.
- Laudan, L. (1984). *Science and values*. Berkeley: University of California Press.
- Laudan, L. (1996). *Beyond positivism and relativism: Theory, method, and evidence*. Boulder: Westview Press.
- Marshall, B. (2002). The discovery that *Helicobacter pylori*, a spiral bacterium, caused peptic ulcer disease. In B. Marshall (Ed.), *Helicobacter pioneers: First hand accounts from the scientists who discovered Helicobacters* (pp. 165–202). Victoria, Australia: Blackwell Science Asia.
- Palmer, E. (1954). Investigations of the gastric mucosa spirochetes of the human. *Gastroenterology*, 27, 218–220.
- Peterson, W. L., Harford, W., & Marshall, B. J. (2002). The Dallas experience with acute *Helicobacter pylori* infection. In B. Marshall (Ed.), *Helicobacter pioneers: Firsthand accounts from the scientists who discovered Helicobacters* (pp. 143–150). Victoria, Australia: Blackwell Science Asia.
- Popper, K. (1975). The rationality of scientific revolutions. In R. Harre (Ed.), *Problems of scientific revolution: Progress and obstacles to progress*. Oxford: Clarendon Press.
- Rigas, B., & Papavassiliou, E. D. (2002). John Lykoudis: The general practitioner in Greece who in 1958 discovered the etiology of, and a treatment for, peptic ulcer disease. In B. Marshall (Ed.), *Helicobacter pioneers* (pp. 75–87). Oxford: Blackwell publishers.
- Robbins, H. (1952). Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society*, 58, 527–535.
- Sarkar, H. (1983). *A theory of method*. Berkeley: University of California Press.
- Solomon, M. (1992). Scientific rationality and human reasoning. *Philosophy of Science*, 59(3), 439–455.
- Solomon, M. (2001). *Social empiricism*. Cambridge, MA: MIT Press.
- Strevens, M. (2003a). Further properties of the priority rule. Manuscript.
- Strevens, M. (2003b). The role of the priority rule in science. *Journal of Philosophy*, 100(2), 55–79.
- Thagard, P. (1993). Societies of minds: Science as distributed computing. *Studies in History and Philosophy of Science*, 24, 49–67.
- Thagard, P. (1998a). Ulcers and bacteria I: Discovery and acceptance. *Studies in History and Philosophy of Science. Part C: Studies in the History and Philosophy of Biology and Biomedical Sciences*, 29(1), 107–136.

- Thagard, P. (1998b). Ulcers and bacteria II: Instruments, experiments and social interactions. *Studies in History and Philosophy of Science. Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 29, 317–342.
- Unge, P. (2002). *Helicobacter pylori* treatment in the past and in the 21st century. In B. Marshall (Ed.), *Helicobacter pioneers: Firsthand accounts from the scientists who discovered Helicobacters* (pp. 203–213). Victoria, Australia: Blackwell Science Asia.
- Warren, J. R., & Marshall B. J. (1984). Unidentified curved bacilli on gastric epithelium in active chronic gastritis. *Lancet*, 1(8390), 1311–1315.
- Weisberg, M., & Muldoon R. (2008). Epistemic landscapes and the division of cognitive labor. *Philosophy of Science*, forthcoming.
- Zollman, K. J. (2007). The communication structure of epistemic communities. *Philosophy of Science*, 74(5), 574–587.