

Social structure and the effects of conformity

Kevin James Spears Zollman

Received: 12 March 2008 / Accepted: 20 August 2008
© Springer Science+Business Media B.V. 2008

Abstract Conformity is an often criticized feature of human belief formation. Although generally regarded as a negative influence on reliability, it has not been widely studied. This paper attempts to determine the epistemic effects of conformity by analyzing a mathematical model of this behavior. In addition to investigating the effect of conformity on the reliability of individuals and groups, this paper attempts to determine the optimal structure for conformity. That is, supposing that conformity is inevitable, what is the best way for conformity effects to occur? The paper finds that in some contexts conformity effects are reliability inducing and, more surprisingly even when it is counterproductive, not all methods for reducing its effect are helpful. These conclusions contribute to a larger discussion in social epistemology regarding the effect of social behavior on individual reliability.

Keywords Conformity · Social network · Social structure · Social epistemology · Agent based model

1 Introduction

In recent years there has been increasing interest in the relationship between social institutions and reliable inference. So called social epistemology focuses on the effect of social behavior on the reliability of individuals or groups (e.g. Goldman 1999). There are two broad strategies employed in approaching this problem. One might presume that individuals are behaving so as to maximize their reliability and ask what social institutions and structures would render them most effective. The construction of these “scientific utopias” fails to offer normative recommendations to

K. J. S. Zollman (✉)
Carnegie Mellon University, Pittsburg, PA, USA
e-mail: kzollman@andrew.cmu.edu

groups composed of actual individuals since many individuals fail to achieve optimal epistemic behavior.¹ Alternatively, one might make reasonable assumptions about the behavior of individuals and ask what social structures will make them more or less reliable.² This project allows for the comparison of actual social institutions and also enables us to defend the value of contemporary social institutions against critics.³ This paper contributes to this later project.

Our focus here will be on one, oft criticized, feature of human belief formation: social conformity. In investigating this behavior we endeavor to answer two questions. First, is conformity epistemically good, i.e. does it increase the reliability of the group or individual? Second, given that people conform, what social structure maximizes the reliability of the group and of the individuals? With respect to this later question we will focus on the structure of interactions. What is the optimal distribution of influence? These questions will be answered by proposing a mathematical model of social conformity in situations where individuals are judging from two available options (e.g. 'true' and 'false'). We conclude from this model that social conformity can, in some cases, be reliability inducing.

In answering the second question, we will compare the effect of conformity in several different social networks. These networks are taken to model the structure of social interaction, and the conclusions might offer some recommendation for the optimal social structure of epistemic groups. Here we find that the optimal social structure is one where individuals are influenced by the largest number possible. This suggests that one ought not be interested in minimizing the degree of social influence if it cannot be totally eliminated, since this would have a counterproductive effect.

2 Social psychology and conformity effects

It is well known that humans will readily conform to the wishes or beliefs of others. It was perhaps a surprise when Solomon Asch (1955, 1956) found that people will do this even in cases where they can obviously determine that others are incorrect. Asch presented subjects with two cards, one contained a single reference line and the other contained three lines of various lengths (one was the same length as the reference line). Asch manipulated the social situation by occasionally having two confederates publicly answer incorrectly prior to the subject providing an answer. The subject heard the incorrect responses of the others and was asked to publicly declare his answer as well. Asch found that the degree of conformity was relatively high.

It has been suggested that these experimental subjects offered incorrect answers because they viewed answering incorrectly as being relatively low cost. They might have evaluated the cost of disagreement as higher than the cost of being incorrect,

¹ The construction of these utopias is undertaken, and defended, by Starbuck (2007).

² I take this to be the project of many recent investigations in social epistemology by Heugelmann and Krause (2002, 2006), Kitcher (1990, 1993, 2002) and Strevens (2003, 2006).

³ For instance, one project by Kitcher (1993) is to demonstrate that apparently counterproductive behavior in scientists is actually reliability inducing. One target of this argument are critics who point to these behaviors as evidence of the non-truth-tracking nature of science.

⁴ Earlier work on optimal network structure has been done by Eyalman (2007a, b).

and thus decided that lying (in this context) was the better option. There is qualified support for this claim. A later experiment slightly modified Asch's study by priming the subjects differently (Baron et al. 1996). One group was told the study was largely irrelevant and then asked to complete the task. Another group was told the study was very important and paid for correct answers. In the later group conformity was markedly decreased. However, when the same setup was repeated with a more difficult task, conformity *increased* as subjects perceived the outcome to be more important! This suggests that when confronted with difficult situations we may use others as informational sources when attempting to determine facts about the world (Baron et al. 1996⁵).

Conformist behavior is often criticized as representing a shortcoming of human belief formation. About this Asch said,

Consensus, to be productive, requires that each individual contribute independently out of his experience and insight. When consensus comes under the dominance of conformity, the social process is polluted and the individual at the same time surrenders the powers on which his functioning as a feeling and thinking being depends. That we have found the tendency to conformity in our society so strong that reasonably intelligent and well-meaning young people are willing to call white black is a matter of concern. It raises questions about our ways of education and about the values that guide our conduct (Asch (1955 p. 34).

While concern over intentional manipulation is serious, it is not at all clear if conformist behavior is inappropriate in a context where one imagines oneself in a group of equally reliable and honest peers (as the subjects of these experiments were carefully misled to believe). In such a case, when confronted with a majority of peers who disagree it might be that one ought to believe that oneself is mistaken and one's peers are correct.⁶

This simple argument cannot provide a compelling reason to regard conformity as inducing reliability. Ignoring the possibility of being manipulated, one often does not have direct access to other individuals' perceptions of the world. Their current reports may have been influenced by other individuals; two individuals may well come to believe something because they were influenced by the same person. Since their judgments are no longer independent, one does not have the same reason to adopt the majority view. Because of concerns of this sort, we must complete a more complex analysis of the effect of conformity on the reliability of individuals and groups if we are to determine more exactly the epistemic effects of conformity.⁷

⁵ It should be noted that the "more difficult" task was still not particularly difficult. A control to determine the difficulty found that without the biasing effect of other answers, subjects answered correctly 75% of the time. Since in this experiment there were four available answers, these subjects were three times more likely than chance to answer correctly.

⁶ Thinking in terms of reliability, if one's peers are equally reliable, better than chance, and independent one has a higher probability of believing the truth if one adopts the majority view. Insofar as one ought to do what makes one more reliable, in this situation, one ought to conform (Bertrjee 1992).

⁷ Studying the epistemic effects of conformity is not new with this study. In addition to the predecessors noted in the next sections, others have developed other models which study the effect of conformity like effects. See Brock and Durlauf (2002), Hegselmann and Kraus (2006), Weisberg and Muldoon (forthcoming), and DeMarzo et al (2003).

3 A model of conformity in deliberation

In order to investigate the effects of conformity on epistemic reliability, we will study one simple model of conformity effects. While this model is impoverished in many ways relative to the actual effect of conformity on individuals, its simplicity will provide one avenue for investigating the effects of the phenomenon.

Suppose there are two beliefs. We will presume that one is true, named a , and another is false, named b . Each individual in our model will be informed (with noise) about this fact. With probability $1 - \epsilon$ the individual is correctly informed and with probability ϵ she is told a falsehood. Individuals are arranged on a graph and time is divided into a series of discrete time periods. In each period every individual simultaneously surveys her neighbors' beliefs, if the majority of people she sees (including herself) believes the opposite, the individual changes her mind. Otherwise she keeps it the same.

Formally, we will represent an individual i 's belief at a time t by S_i^t . We will represent i 's neighborhood (those to whom she is connected, including herself) as N_i . Let $N_i^t(T) = |\{a \in N_i : S_a^t = T\}|$, i.e. the number of individuals in i 's neighborhood who believe T at time t . $N_i^t(F)$ can be defined similarly. With this in hand we can represent an individual's belief at a time as a function of her neighbors' beliefs on the prior time step.

$$S_i^t = \begin{cases} T & \text{If } N_i^{t-1}(T) > N_i^{t-1}(F) \\ T & \text{If } N_i^{t-1}(T) = N_i^{t-1}(F) \text{ and } S_i^{t-1} = T \\ F & \text{Otherwise} \end{cases} \quad (1)$$

Consider a graph G , which is composed of a set of nodes V and a set of edges E . We will define a belief graph as a pair (G, f) where G is the graph and f is a function from individual to $\{T, F\}$ (representing an individual's belief). Let the set of belief graphs be B , using Eq. 1 we can define a function $B \rightarrow B$ which represents the change of the system during one time step. We will call a system converged if it represents a fixed point of f , i.e. ab such that $f(b) = b$. As will be seen below, for finite groups, it need not ever converge.

This model can have many possible end states. For our purposes, we will divide them into four categories.

1. *Correct unanimity* The system converges to a state where everyone believes
2. *Correct majority* The system converges to a state where a majority of the individuals believe T .

⁸ While inclusion of an individual in her own neighborhood may seem an innocuous assumption (Poljak and Sura 1983), (1998) found some significant differences in models which include this assumption. Given our purposes, it seems appropriate that the individual consider her own belief, but it is worth noting that some of these results might be different if she did not.

⁹ Poljak and Sura (1983) and Goles and Olivos (1980) both independently proved that such a system has either period one or two. That is, it either converges to a fixed point, where it remains, or converges to a state which repeats every second stage.

3. *Incorrect convergence* The system converges to a state which is neither *correct unanimity* or *correct majority*.
4. *Non-convergence* The system never converges.

This model, or very similar models, have already employed in a wide variety of fields, but interest in these models has primarily been for other purposes.¹⁰ Most investigations have focused on three things: the ability of a limited number of individuals to influence the final outcome, the resilience of this model to limited changes, and the ability of this local learning rule to compute global properties of the system. While a few of these results can be employed here, further study is needed in order to answer our questions.

This model also bears a strong similarity to ‘information cascade’ models used in economics (Banerjee 1992; Welch 1992; Bikhchandani et al. 1992). These models focus on economic decisions, where individuals reveal their beliefs through actions (like buying stock). In the information cascade models once an individual reveals her opinion, she can no longer change it and there is a rigid temporal order to individuals’ actions. The model presented in this paper allows individuals to choose and then later change their mind based on the choices of others, which conforms more directly to the process of mutual influence where individuals can affect one another.

4 Conformity in networks

We will restrict our attention to three canonical networks, the cycle, the wheel, and the complete graph (see Fig. 1). These networks provide a nice sampling of the population of graphs. Two of them (the cycle and the complete graph) are very homogeneous, but differ in terms of connectivity. The wheel is less homogeneous, because it has a central individual who is connected to everyone.

These networks are idealizations from three different social circumstances, and while none of them directly represent the actual society in which individuals find themselves, they do provide a helpful starting point. The cycle, represents a situation of perfectly symmetric influence. No one individual is any more influential than any other and all individuals have some (indirect) influence over everyone else. Here we might imagine a group of individuals which rarely interact with others, and no one person interacts with everyone. The complete graph is also perfectly symmetric, but all influence is direct. In the complete graph represents a circumstance where everyone knows everyone else—a tight knit group. The wheel represents a circumstance where the symmetry of influence is broken. For example, a loosely associated group with a single, influential person. With these three networks we can consider the effect of both direct versus indirect influence and symmetric versus asymmetric influence.

In order to analyze the epistemic properties of different groups we must be explicit about what constitutes success. For the individual, this is easy. If an individual believes

¹⁰ This model has been used in mathematics (Dorsett and Steif 1993; Goles and Olivos 1980; Granville 1991; Holley and Liggett 1975; Moran 1994; ϕ , 1995; Mustafa and Pekev 2004; Steif 1994), computer science (Flocchini et al. 1998; Hassin and Peleg 2000; Kržišljović 2001; Nakata et al. 1999; 2000; Peleg 1998; 2002), biology (Agur 1987; Agur et al. 1988; Clifford and Sudbury 1973) and social psychology (Latané and Nowak 1997; Poljak and Sura 1993).




Fig. 1 A 10 person cycle, wheel, and complete graph

T , he has succeeded. For the group, it may be more difficult. One natural metric to measure how often the group achieves unanimity on the correct answer. As will be seen, unanimity is sometimes difficult to come by in this model. We will, in addition to correct unanimity, use correct majority as a metric of success. Correct majority probably represents the weakest measure of success and as such provides a sort of minimal condition for our consideration.

With this in hand we will now look at the effect of the three canonical networks on our three measures of success (one individual, two group).

4.1 Complete networks

In the case of the complete network the two group features are identical and the individual one is very similar. Since everyone sees everyone else, there are only three possible convergent states. One where everyone believes T , one where everyone believes F , and one where exactly half believe each.

Calculating the probability of each of these involves only a simple application of the binomial theorem.

$$U_{COM}(n, \epsilon) = \begin{cases} \sum_{i=(n/2)+1}^n \binom{n}{i} (1-\epsilon)^i \epsilon^{n-i} & \text{if } n \text{ is even} \\ \sum_{i=(n+1)/2}^n \binom{n}{i} (1-\epsilon)^i \epsilon^{n-i} & \text{otherwise} \end{cases} \quad (2)$$

This represents the probability that the group converges to correct unanimity (or majority). In order to calculate the probability that an individual in this group is correct, one need only add $(1/2)\epsilon^{(n/2)}(1-\epsilon)^{(n/2)}$ to the even case.

The analysis of the complete network is identical to the Condorcet Jury Theorem. While the interpretation of the model is different in the Condorcet Jury Theorem individuals have a belief and then vote on an outcome where the majority wins the mathematical facts are the same.

In the complete network only even numbered networks which initially have one half believing T and one half believing F will not converge to a state of unanimity. This demonstrates a fact already known about the Condorcet Jury Theorem, that increasing the number in the jury does not monotonically increase the reliability of that jury. Instead, a weaker statement must be proven.

¹¹ Ben-Yashar and Parous (2000) attribute the first proof of this to Laplace.




Fig. 2 Complete networks for 9 error probabilities

Proposition 1 For any $\epsilon < 0.5$, for any n , $U_{COM}(n, \epsilon) < U_{COM}(n + 2, \epsilon)$. Additionally, for any $\epsilon < 0.5$, for any even n , $U_{COM}(n, \epsilon) < U_{COM}(n + 1, \epsilon)$.

That is, for groups of reliable individuals, adding two members to any group or one member to an even sized group makes that group more reliable. For groups made up of unreliable individuals the opposite is true, as they grow they become less reliable.

Proposition 2 For any $\epsilon > 0.5$, for any n , $U_{COM}(n, \epsilon) > U_{COM}(n + 2, \epsilon)$. Additionally, for any $\epsilon > 0.5$, for any odd n , $U_{COM}(n, \epsilon) > U_{COM}(n + 1, \epsilon)$.

For visualization, Eq. 2 is graphically represented in Fig. 2. Here it is easy to see both Propositions 1 and 2. Those whose probability of error is below 0.5 are trending toward 1.0 as the size grows and those whose probability of error is above 0.5 are trending toward 0.

Proposition 1 has exactly the same interpretation here as is used in the Condorcet Jury Theorem, that groups of individuals are more reliable than any single member of the group (taken alone) when the individuals are better than chance. In this sense, individuals who have access to everyone in the group do better by conforming than ignoring the others.

Equation 2 not only represents the complete graph. For a slightly different dynamic,¹² several theorems from Mustafa and Peke (2004) entail that many other graphs have the same convergence function. All of these networks are very highly connected and are thus very close to the complete graph. Mustafa and Pekev show that any graph with this convergence function has the property that everyone's second neighborhood (their neighbors' neighbors) must be at least a majority of the graph.

¹² The only difference is that individuals adopt what the majority of their neighbors believe, and staying the same in the case of a tie.

They also conjecture that every graph with this convergence function will include at least one person who is connected to everyone.¹³

4.2 Cycle

The connected graph represents a situation where everyone is equally influenced by everyone else. In less connected societies, people may not have access to the beliefs of everyone in the community. Instead they may only be acquainted with the beliefs of a few other individuals. In order to analyze a situation like this we will turn to a graph which represents another informational extreme, the cycle.

Unlike the complete graph, in the cycle many different outcomes are possible. Since each individual only sees three people (herself and her two neighbors), an individual will only switch her belief from one period to the next if both of her neighbors hold a belief opposite her own. This means that any two neighboring believers will remain that way forever. This fact significantly increases the number of fixed points, but we are only interested in two classes of final outcomes: the outcome where everyone believes T and the set of outcomes where the majority believe F . We will analyze these in turn.

4.2.1 Unanimity

Because any two neighboring believers will never change their minds, if a network begins with two neighboring F believers, then the network will never converge to correct unanimity. In addition, a cycle where every believer is surrounded by F believers and vice versa will never converge but will instead oscillate with every individual changing her mind every round. Happily, these two initial states represent all the initial states which do not converge to correct unanimity. This is represented by the following proposition, whose proof is in the Appendix.

Proposition 3 *A cycle belief graph will converge to correct unanimity if and only if (a) it has two neighboring T believers and (b) it does not have two neighboring F believers.*

With this in hand, we are now able to calculate the probability that a random initial starting configuration satisfies this condition. In order to do this, we will consider a simpler situation, analogous to a coin flipping problem. Consider a string of flips of a coin that has probability $(1 - \epsilon)$ for heads (or T). We must calculate the probability that this string does not have two tails (or F s) in a row. We will define this function, Q recursively.

Since for strings of length one it is guaranteed that the string does not contain two tails, $Q(1, \epsilon) = 1$. One can easily calculate that $Q(2, \epsilon) = 2(1 - \epsilon) - (1 - \epsilon)^2$. For the recursive step, we will consider two cases. Suppose the coin is flipped. If the first flip is heads, the probability the coin does not have two tails in a row is

¹³ These two conditions are not sufficient, however, since the wheel satisfies both conditions but has a very different function.

$Q(n - 1, \epsilon)$ (this equals the probability that the remaining $n - 1$ bits do not contain two tails). If the first bit is tails, the probability that the coin does not have two tails in a row is $(1 - \epsilon)Q(n - 2, \epsilon)$ (the probability the second bit is heads and then that the remaining string does not have two tails). So we have:

$$Q(1, \epsilon) = 1 \tag{3}$$

$$Q(2, \epsilon) = 2(1 - \epsilon) - (1 - \epsilon)^2 \tag{4}$$

$$Q(n, \epsilon) = (1 - \epsilon)Q(n - 1, \epsilon) + \epsilon(1 - \epsilon)Q(n - 2, \epsilon) \tag{5}$$

Because we are only requiring two tails, regardless of the size of the run, as the length of the string increases, the probability that this substring appears increases. As a result, we can easily prove that it is decreasing.

Proposition 4 For $0 < \epsilon < 1$, $Q(\cdot, \epsilon)$ is strictly decreasing.

Proof Since $\epsilon < 1$, $Q(1, \epsilon) = 1 > Q(2, \epsilon)$. Suppose that there is some n such that for all $x < n$ and all $y, x < y \leq n$, $Q(x, \epsilon) > Q(y, \epsilon)$, to prove that $Q(n, \epsilon) > Q(n + 1, \epsilon)$

$$\begin{aligned} Q(n, \epsilon) - Q(n + 1, \epsilon) &= (1 - \epsilon)Q(n - 1, \epsilon) + \epsilon(1 - \epsilon)Q(n - 2, \epsilon) \\ &\quad - (1 - \epsilon)Q(n, \epsilon) - \epsilon(1 - \epsilon)Q(n - 1, \epsilon) \\ &= (1 - \epsilon)(Q(n - 1, \epsilon) - Q(n, \epsilon)) \\ &\quad + \epsilon(1 - \epsilon)(Q(n - 2, \epsilon) - Q(n - 1, \epsilon)) \end{aligned}$$

By the inductive hypothesis $Q(n - 1, \epsilon) - Q(n, \epsilon) > 0$ and $Q(n - 2, \epsilon) - Q(n - 1, \epsilon) > 0$. Since $0 < \epsilon < 1$, $Q(n, \epsilon) - Q(n + 1, \epsilon)$ is positive. \square

In this idealization we have ignored one connection, the cycle contains one additional set of neighboring individuals (individual 1 is connected to individual n). For $n > 2$, we can now define a function P , based on Q . P gives the probability that no two neighbors believe F for a given size n . We will construct P by considering a cycle and arbitrarily breaking one edge. This provides us with a string of individuals, which we can enumerate $1, \dots, n$.

We now calculate the probability that this string does not contain two consecutive F believers and that the first and last individual in the string do not believe F . This is done, as before, by breaking the situation into two parts. When the first individual believes T , the question is only if the remaining string of $n - 1$ individuals contains two neighboring F believers. If the first individual believes F , then the network does not have two consecutive F believers if and only if the second and last individuals in the string believe T and the remaining string of $n - 3$ individuals do not contain two consecutive F believers. Again, we can define the function, this time in terms of

$$P(3, \epsilon) = 3(1 - \epsilon)^2 - 2(1 - \epsilon)^3 \tag{6}$$

$$P(n, \epsilon) = (1 - \epsilon)Q(n - 1, \epsilon) + \epsilon(1 - \epsilon)^2Q(n - 3, \epsilon) \tag{7}$$

It immediately follows from Proposition 4 that,

Proposition 5 For $0 < \epsilon < 1$, $P(\cdot, \epsilon)$ is strictly decreasing.

For even n , P includes networks which alternate between C and F (which do not converge). We must discount this probability for even networks. With that discounting we have the following correct unanimity convergence function for the cycle:

$$U_{CYC}(n, \epsilon) = \begin{cases} P(n, \epsilon) - 2\epsilon^{(n/2)}(1 - \epsilon)^{(n/2)} & \text{if } n \text{ is even} \\ P(n, \epsilon) & \text{otherwise} \end{cases} \quad (8)$$

Because of this extra term, U_{CYC} is not decreasing in n . It appears that for large n , U_{CYC} is decreasing, which can be easily seen in Fig. 3 (the graphical representation of Eq. 8). As the network size grows the probability that chance will result in two neighboring individuals both believing C goes up. Since this is all that is required for a cycle to fail to converge to correct unanimity, the probability that a cycle correctly converges to correct unanimity goes down as the network grows. Although it appears that U_{CYC} continues to monotonically decrease as n grows (for given ϵ), the complexity of the function makes proving this fact rather difficult. We can, however, easily see that $\lim_{n \rightarrow \infty} U_{CYC}(n, \epsilon) = 0$.

4.2.2 Majority

We now turn to analyzing the probability that a cycle converges to at least a simple majority believing the correct belief. Unfortunately, computing functions which represent this probability is sufficiently complex that analytical treatment is no longer possible. However, the state space is relatively small. Instead of deriving a function or simulating, we will use computational methods to completely search the state space for a given n and ϵ . The results reported below have exhausted the possibilities for the given values, the only possibility for error is in programming error. Unfortunately, however, the results are not sufficiently general to prove facts about thresholds.




Fig. 3 The cycle for 9 error probabilities (correct unanimity)




Fig. 4 The cycle for 9 error probabilities (correct majority)

The probabilities for selected \tilde{O} s and \tilde{O} s are pictured in Fig. 4. With this more weak judgment of success in hand, we find that the cycle behaves more like the complete graph. It appears that as two individuals are added those groups composed of reliable individuals become more reliable, while groups made up of unreliable individuals become less reliable. Unlike the complete graph, one does not need reasonably large groups before the group becomes more reliable than the individual. For example when $\epsilon = 0.4$, one needs only seven members before the group is more reliable than any of its individuals (with reliability judged by convergence to correct majority).

Here we see the first example of a situation where the judgment of success makes a difference in the comparison of groups. Judged by convergence to correct unanimity we would regard smaller cycles as superior to larger ones. But, judged by convergence to correct majority, we would regard larger cycles as superior. This illustrates that the choice of a group metric for reliability can be very important.

4.2.3 Individual

We began with a question about the epistemic benefit of conformity effects on the reliability of individuals. In order to evaluate the individualistic effects of conformity we will compare the number of individuals who believe in the enumeration used above to ϵ , the measure of individual reliability. Figure 6 shows the relative increase or decrease of an individual's reliability from participating in the conformity network. For $\epsilon < 0.5$ (where individuals are more reliable than chance), we find that the social interactions improve their reliability. For $\epsilon > 0.5$ we find that participation makes them less reliable than they would be on their own.

There are two interesting features of these results that are worth noting. First is the relationship between ϵ and the difference in reliability. When individuals are very reliable ($\epsilon = 0.1$) the benefit of participating in the group is less than when they are slightly less reliable ($\epsilon = 0.2$). This suggests that conformity has a larger effect on individual reliability in those cases where the individual is somewhat reliable, but not extremely




Fig. 5 Comparing the cycle and individual reliability

so. It is interesting to notice that the point at which the conformity effect results in the largest growth of reliability (around $\epsilon = 0.2$ to 0.3) is exactly the difficulty of the difficult task in [Baron et al. \(1996\)](#). Recall that in the difficult task, people tended to increase their propensity to conform. While it would be reckless to presume a causal relationship between these facts, it does provide a nice justification for such behavior.

Second, although we are only looking at the reliability of single individuals here, we find that group properties are reflected in individual reliability. In particular, group size has an important effect. This is somewhat surprising since, regardless of the network size, the individuals are only connected to two other people. However, the size of the network does have an influence of the reliability of individuals, with larger groups resulting in higher reliability.

4.3 Wheel

The final network that we will consider is the wheel. The wheel is identical to the cycle, except one of the individuals is connected to every other one. This network is often used to represent a particularly influential or central person who communicates with everyone (while many of them do not communicate with each other). Here we can judge the effect of unequal influence against the previous two homogeneous graphs. Again, we will begin our discussion with convergence to correct unanimity.

4.3.1 Unanimity

Each individual on the outside (i.e. not the central individual) has three neighbors. Since he includes himself and only changes if a majority believe the opposite, an individual on the outside will only change if all of his neighbors believe the opposite. As a result, if any two individuals on the outside of the wheel believe the network will not converge to correct unanimity. We can again prove the conditions under which the network correctly converges.

Proposition 6 *The wheel will not correctly converge if and only if (a) two individuals on the outside both believe F or (b) there is an individual (not the hub) all of whose neighbors believe F .*

Proof (i) If a network satisfies either (a) or (b) it will not converge. As was the case with the cycle, if any of the exterior agents have at least one neighbor that agrees they will not change their beliefs. Therefore, a network satisfying (a) will not converge to correct unanimity. Similarly for (b) our individual's neighbors (not the hub) see the hub who believes F and will not change their beliefs. Our focal individual will believe F , however, since all three of his neighbors believe F . This produces a chain of FF on the exterior which prevents the network from converging to correct unanimity.

(ii) If a network does not satisfy (a) or (b) it will correctly converge. Suppose β does not satisfy (a) or (b). Then β does not have βF and either (b1) the hub believes T or (b2) there are at least two T 's separating β and the hub.

Suppose β fails to satisfy (a) and satisfies (b1). Because β fails to satisfy (a), every F believer will change because all of his neighbors believe T . The T believers see the center, who believes T , and will not change. Because β fails to satisfy (a), at most half of the outside can believe F and as a result the center continues to believe T . After one round everyone believes T .

Suppose β fails to satisfy (a) and satisfies (b2). Since β 's neighbor is surrounded by F believers β will change. Since β satisfies (b2), there must be twice as many T 's as F 's. Since we are considering networks of size 5 or larger (with outsides of size 4 or larger), this results in the center changing. On the next time step β 's neighbor is surrounded by T believers; they will all switch while the center remains the same. After this the network has converged to correct unanimity. \square

In order to develop a function that expresses the probability of convergence to correct unanimity for this network, we must consider two cases. In the case where the center believes T (which occurs with probability $1 - \epsilon$), unanimity is achieved with probability $P(n - 1, \epsilon)$ (for P as defined in the preceding section). If the center believes F , unanimity is achieved if the outside does not contain a string of FF or FTF . This calculation is more difficult.

In order for the outside to not contain any string of this type it must have at least two neighboring individuals (both on the outside) who believe T . This occurs with probability $(1 - P(n - 1, (1 - \epsilon)))$. We can then take this group and break the connection between those two individuals to form a string which begins and ends with T . We can then treat the intervening individuals as a Markov chain with four states. The states are: (1) Next two must be TT , (2) Next one must be T , (3) Next can be either T or F , and (4) Failed. Where "Failed" represents a network which does contain the string FF or FTF . The transition probabilities for this chain are given by:

$$\mathbf{M} = \begin{pmatrix} 0 & (1 - \epsilon) & 0 & \epsilon \\ 0 & 0 & (1 - \epsilon) & \epsilon \\ \epsilon & 0 & (1 - \epsilon) & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \tag{9}$$




Fig. 6 The wheel for 9 error probabilities (correct unanimity)

With this in hand we can define the convergence function:

$$U_{WHE}(n, \epsilon) = (1 - \epsilon)P(n - 1, \epsilon) + \epsilon(1 - P(n - 1, (1 - \epsilon)))(1 - M_{34}^{(n-3)}) \quad (10)$$

The reader will undoubtedly be unaffected by this function, a graphical representation of it is presented in Fig. 6. Here we see that, like the cycle, for all more individuals make the wheel less reliable. Like the cycle, we can demonstrate that $U_{WHE}(n, \epsilon)$ approaches 0 as $n \rightarrow \infty$. As we have already shown in Proposition 6, for any ϵ , $\lim_{n \rightarrow \infty} P(n, \epsilon) = 0$. Similarly, because M has a unique absorbing state which is accessible from any state, $\lim_{n \rightarrow \infty} M_{34}^n = 1$. Since $P(n, \epsilon)$ is decreasing in n and has a limit of 0, $\lim_{n \rightarrow \infty} (1 - P(n, 1 - \epsilon)) = 1$, and, since $P(\cdot)$ is probability, it is bounded above by 1 and below by 0. As a result, $\lim_{n \rightarrow \infty} \epsilon(1 - P(n - 1, 1 - \epsilon))(1 - M_{34}^{n-3}) = 0$.

4.3.2 Majority

Again we will use computer enumeration to determine the probability that the wheel converges to correct majority. The results are pictured in Fig. 7. Again we find the same increasing structure present in both the complete graph and cycle (as judged by convergence to correct majority). Judged by this metric, it again appears that conformity has made our group more reliable.

4.3.3 Individual

The results for the individuals in the wheel are qualitatively identical to the results obtained for the individuals in the cycle (pictured in Fig. 8). We find that the influence of the network is greatest when n is small but not as small as possible. Also, we again find that the size of the network projects down onto individual reliability.




Fig. 7 The wheel for 9 error probabilities (correct majority)

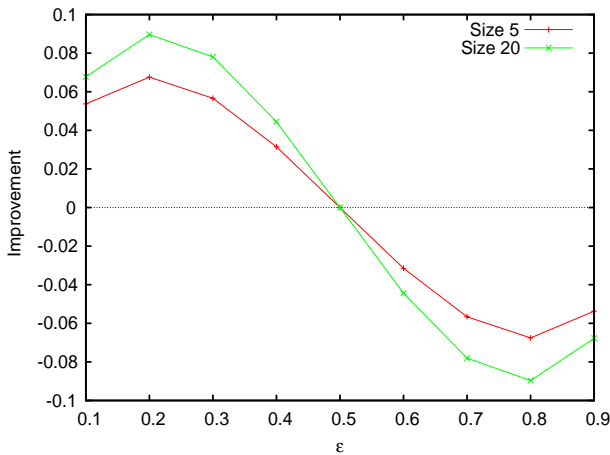


Fig. 8 Comparing the wheel and individual reliability

4.4 Comparisons

So far we have primarily focused on the benefit or harm of conformity on the reliability of three different groups. We have considered reliability from three different perspectives: group unanimity, group majority, and individual reliability. A second question troubled us at the outset. Given that individuals conform, what is the optimal network structure (in terms of reliability)? In this section we will compare the three networks already analyzed and compare their reliability using the same three metrics employed above.

4.4.1 Unanimity

For $\epsilon < 0.5$, the complete graph trends upwards, but the cycle and wheel trend down. In these cases it is clear that for large n the complete graph will be superior. Analytical results beyond this are difficult to come by, however, Fig. 9 shows the three graphs for two different values of ϵ .

This graph suggests that for sufficiently large n (in this case, not very large), the complete graph is superior to the wheel and the cycle. The relationship between the wheel and the cycle here is interesting. For small n the wheel is superior, but for larger n the cycle is superior. While both functions approach 0 in the limit, it appears that the wheel does so faster.

4.4.2 Majority

When we compare the three networks using the majority success function the similarity between the three is striking (see Fig. 10). Here all three networks almost completely overlap. Interestingly, there is no fact about connectivity that is true for all $\epsilon = 0.3$, the complete network is superior to the cycle for odd networks and only slightly inferior for even sized ones. For $\epsilon = 0.7$, however, the cycle is superior to the complete network almost always. The cycle appears to be effected by even/odd switches less substantially than the complete network, but it otherwise runs very close to this other graph.

Both of these networks appear to be better than the wheel when agents are more reliable than chance and worse when they are less. This suggest that the unequal connectivity of the wheel has a negative effect on the overall reliability of groups composed of reliable members.

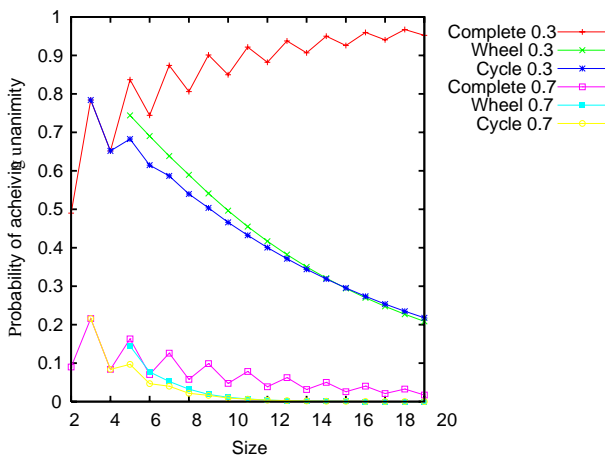


Fig. 9 A comparison for two values of ϵ




Fig. 10 Comparison for a majority success rule

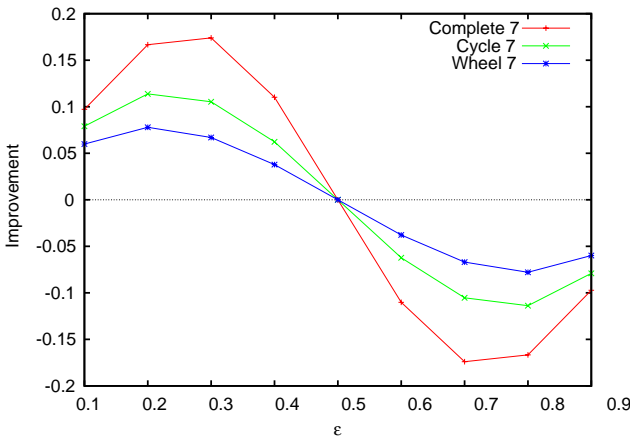


Fig. 11 Comparing the three networks and individual reliability

4.4.3 Individual

We can also compare the reliability of the individuals in the group across the three networks (see Fig.1). Here we find exactly the same results as the other measures of reliability. When the individuals are better than chance, the complete network is superior to the cycle which is superior, in turn, to the wheel. The relationship reverses when the individuals become unreliable. For all three networks we observe the same underlying relationship, that the network helps the most when the individuals are somewhat reliable but not incredibly so.

4.5 Mutation

Many of the results of the previous sections depended on the deterministic character of the model. For example, the ‘blinking’ state where individuals on the cycle change their beliefs on every round will disappear if we add even the slightest chance that an individual will fail to change from one round to the next. This feature makes this result unstable to very small changes in the rules by which players update, and thus not terribly helpful in predicting or explaining actual human populations.

In order to determine how serious this effect is, we will consider a model where there is some probability m , on every round that an individual changes her belief to one not specified by the dynamics above. That is, an individual will stay the same when she should switch or switches when she should stay the same.

Consider a cycle where there are only two believers. We will represent their neighborhood thusly,

$$\begin{matrix} \dots T F F T \dots \\ \dots 1 2 3 4 \dots \end{matrix}$$

This network will converge to unanimity if and only if, (1a) individual 2 switches to T before individual 4 switches to F or (1b) individual 3 switches to F before individual 1 switches to T , and (2) this happens before two neighboring individuals (other than 2 and 3) switch to F on the same round. For sufficiently small m this will occur reasonably often making it possible for some cycles to converge to correct unanimity in the modified dynamics that would not have in the deterministic one.

However, if m becomes too large (relative to the size of the network) the probability that two neighboring individuals switch becomes large. In this case it will make it more difficult for unanimity to be achieved because too many people are switching. While calculating any of these probabilities is painstakingly complex, we can use computer simulation to estimate them. The results for one such simulation on the 7 person cycle, wheel, and complete graph, for $m = 0.3$, are represented in Table 1.

An important observation about these results is that the move from no mutation to some mutation provides substantial jump in reliability (judged as correct unanimity) for both the wheel and cycle. A mutation probability of only 1 in 10,000 is sufficient to increase both their chances of converging to unanimity by over 10%. Interestingly the increase in mutation rate also has the effect of reducing the disparity between the two different networks. It does not, however, change the relationship between these three, the more connected ones are more reliable. As predicted, more mutation is not always

Table 1 Probability of correct unanimity for $m = 0.3$

Type of network	Mutation probabilities			
	0	0.0001	0.001	0.01
Cycle	0.5866	0.7356	0.8165	0.6512
Wheel	0.6384	0.7792	0.8381	0.6942
Complete	0.8740	0.8740	0.8740	0.8106

Table 2 Probability of correct majority for $\epsilon = 0.3$

Type of network	Mutation probabilities			
	0	0.0001	0.001	0.01
Cycle	08740	08495	08165	06826
Wheel	08387	08369	08380	07923
Complete	08740	08740	08740	08740

Table 3 Probability of individual success for $\epsilon = 0.3$

Type of network	Mutation probabilities			
	0	0.0001	0.001	0.01
Cycle	08054	07964	08246	07521
Wheel	07670	08118	08119	07975
Complete	08740	08740	08733	08740

better. Once a mutation occurs in 1 out of 100 times, the benefit from mutation begins to decrease. Interestingly, even the complete network is harmed.

The results for correct majority are represented in Table 2. Here we find that mutation has much less of an effect and it is almost everywhere negative.

Finally, we can look at the results on individual reliability introduced by mutation. Here we find very little change in overall individual reliability. There are small changes, but the underlying results remain approximately the same. The most substantial change is in the wheel. Mutation here appears to help individuals in the wheel substantially by overcoming the influence of the central figure (Table 3).

Overall the addition of mutation primarily increases the probability of groups converging to correct unanimity, but has little effect on increasing the reliability of the majority or individuals.¹⁵

5 Conclusions

We began our discussion of this model with two questions. The first was, does conformist behavior have a positive effect on the reliability of the individual and of the group? We must provide two answers. When the members are more reliable than chance, an individual does better when he joins a group and engages in conformist behavior than

¹⁴ The reader might find this odd since simultaneous mutation by 4 individuals is required to move a 7 member complete graph from correct unanimity to a state which will converge to incorrect unanimity. This has a 10^{-9} chance of occurring on any given generation. That is not the cause of these results, rather they are the result of fewer individuals mutating on the last round and thus not being eliminated by the dynamics. I leave it to the reader to determine the importance of this result.

¹⁵ Another similar probabilistic model is analyzed by [Lluch-Blassin and Peleg \(2001\)](#) and [Nakata et al \(1999, 2000\)](#). Here individuals adopt the true belief in proportion to the number of true believers in their neighborhood. Such groups always converge to unanimity, but calculating the probabilities of the two states can be quite difficult.

he would have done by merely relying on his own, private judgment. Judged from the individual perspective, one would regard social influence as epistemically productive.

Groups, however, do not necessarily do better when their members participate in a dialog with each other and are affected by conformist behavior. Recall that the reliability of the complete network is identical to the probability of getting the correct answer by taking a majority vote. We discovered that this number was often higher than allowing conformist behavior to affect the beliefs of individuals in other networks and taking a majority vote once those networks had stabilized. This result was most striking with the wheel, where it is clear that from the group perspective we would prefer to take a vote at the beginning rather than allowing individuals to be influenced. Overall it appears that from the group perspective we would prefer the individuals not engage in conformist behavior, but from the individual perspective one would prefer that one does. This result provides an interesting example where the dictates of individual epistemology come apart from the dictates of social epistemology.¹⁶

Our second motivating question was that, given that individuals engage in conformist behavior, what is the optimal social structure for them? Here it appears that the complete graph is most often superior, when individuals are more reliable than chance. This is true both from the individual and group perspective. The complete graph represents a circumstance where social influence is total: individuals are influenced by all the other individuals in a group. The results here suggest that if one has a group of individuals that will be subject to some social pressure, one would like them to be influenced *as much as possible*. Efforts to increase reliability by decreasing, but not eliminating, social influence are counterproductive.

Another important result is that the cycle and the complete graph run close in reliability, while the wheel appears to lag behind when we use majority success. This suggests that unequal connectivity is substantially harmful. This is perhaps the most useful conclusion that can be drawn from these results. Efforts to reduce the effect of social influence are particularly bad when they render the structure of that influence unequal.

Appendix: Proof

We will begin by defining two different types of subgraphs of a given belief graph.

Definition 1 An *a-region* (agreement region) is a subgraph with more than 1 node where all nodes either believe T or F . Let $A(G)$ be the set of all *a-regions* of graph

Definition 2 A *maximal a-region* is an *a-region* which is not a proper subgraph of any other member of $A(G)$. Let $MA(G)$ be the set of all maximal *a-regions* of graph

¹⁶ This is similar to the results from information cascades in economics. Biais et al. (1992) provides an example where information is lost by allowing individuals to observe the choices of others, because they will (rationally) conform to the choices of others and ignore private information. Here too is an example where individuals are behaving in a way that maximizes their individual reliability but the group is worse off as a result.

Informally a maximal a-region is the longest chain of people who all agree. After removing all maximal a-regions, we are left with a (possibly empty) set of regions where all individuals disagree with their neighbors.

Definition 3 A *d-region* (disagreement region) is a subgraph of G such that none of its members are in a subgraph of G . Let $D(G)$ be the set of all d-regions of graph G .

Since we have already removed all the subgraphs with two or more neighbors who agree, d-regions are composed of strings of individuals who have alternating beliefs; every member of a d-region has two neighbors with the opposite belief.

Definition 4 A *maximal d-region* is a d-region which is not a proper subgraph of any other member of $D(G)$. Let $MD(G)$ be the set of all maximal d-regions of graph G .

The underlying strategy is that the graph is partitioned into two types of regions. First, take the largest groups of (at least 2) consecutive F believers, call these maximal a-regions. Second, take the remaining sections and call them maximal d-regions. We can determine the total size of all maximal d- and a-regions $|MA(G)|$ and $|MD(G)|$ be the total number of nodes that are in a maximal a- or d-regions respectively. Notice that $|MA(G)| + |MD(G)| = n$, where n is the total number of nodes.

Lemma 1 Suppose a cycle belief graph G_t and $G_{t+1} = I(G_t)$. $|MA(G_{t+1})| \geq |MA(G_t)|$.

This follows immediately from the fact that individuals will only change their beliefs if both of their neighbors disagree. Since, by definition, all the members of a-regions agree, no member of an a-region will change her mind. As a result the total number of individuals who are members of maximal a-regions cannot shrink.

Lemma 2 Suppose a cycle belief graph G_t where there are both a and d-regions and let $G_{t+1} = D(G_t)$. $|MA(G_{t+1})| > |MA(G_t)|$ and $|MD(G_{t+1})| < |MD(G_t)|$.

Proof We have already demonstrated that MA regions cannot shrink in Lemma 1. First we observe that maximal d-regions must border maximal a-regions. As a result, there are individuals who will change their minds, but will have a neighbor that does not (the individuals who are members of d-regions but who have a neighbor who is a member of an a-region). Since these individuals now share a belief of a neighbor (F), they are now members of a-regions and not members of d-regions. As a result, the number of individuals in a-regions has grown at the expense of individuals in d-regions.

Lemma 3 A cycle belief graph G will converge if and only if there are at least two neighboring individuals who have the same belief.¹⁷

Proof One direction follows immediately from Lemma 2. If a network has two neighboring individuals who believe the same thing, there is an a-region. The a-regions will grow on each round and d-regions shrink until the graph has no d-regions, at which point it has converged. Now it is sufficient to prove that graph that does not have

¹⁷ This fact is observed, without proof, by Agur et al. (1988).

two neighbors who agree will not converge. Suppose such a graph. In this graph every individual disagrees with both of her neighbors. As a result, she will change her belief. This is true regardless of the time, and so as a result, every individual will change her mind on every round.

With one remaining lemma, we can then prove the necessary and sufficient conditions for a cycle network to converge to unanimously believing true.

Lemma 4 *If G_t does not have two neighboring F believers, then $G_{t+1} = I(G_t)$ does not have two neighboring F believers.*

Proof Suppose a network with no two neighboring believers. In the case where the network has no a-regions, this follows immediately from Lemma 3. If the network does have an a-region, that region must be composed of F believers. In the proof of Lemma 2 we demonstrated that the individuals who are in d-regions, but neighbor a member in an a-region will join that a-region on the next round. All those who do not have neighbors in an a-region will change their beliefs on the next round. Since there are no two believers in the d-region, there will be no two neighboring believers in the next round.

As a result we can prove the necessary and sufficient conditions for convergence to correct unanimity.

Proposition 3 *A cycle belief graph will converge to correct unanimity if and only if (a) it has two neighboring T believers and (b) it does not have two neighboring F believers.*

Proof (i) Suppose a graph will converge to correct unanimity. Then we know that it cannot have begun with two F believers, because they constitute an a-region which will not disappear. By Lemma 2 there must have been at least one a-region. As a result it must have been composed of T believers.

(ii) Suppose that a graph has two neighboring T believers and no two neighboring F believers. This means that there is an a-region, and so by Lemma 1 it will converge. By Lemma 4 an a-region of F believers cannot emerge. As a result, the network will converge to all a-regions with no believing a-regions. As a result it has converged to correct unanimity. \square

References

- Agur, Z. (1987). Resilience and variability in pathogens and hosts. *Journal of Mathematics Applied in Medicine and Biology*, 4, 295–307.
- Agur, Z., Fraenkel, A., & Klein, S. (1988). The number of fixed points of the majority rule. *Discrete Mathematics*, 70, 295–302.
- Asch, S. E. (1955). Opinions and social pressure. *Scientific American*, 193(5), 31–35.
- Asch, S. (1956). Studies of independence and submission to group pressure: I. A minority of one against a unanimous majority. *Psychological Monographs*, 70(9).
- Banerjee, A. V. (1992). A simple model of herd behavior. *The Quarterly Journal of Economics*, 107(3), 797–817.
- Baron, R. S., Vandello, J. A., & Brunsman, B. (1996). The forgotten variable in conformity research: Impact of task importance on social influence. *Journal of Personality and Social Psychology*, 71(5), 915–927.
- Ben-Yashar, R., & Paroush, J. (2000). A nonasymptotic condorcet jury theorem. *Journal of Economic Choice and Welfare*, 17, 189–199.
- Bikhchandani, S., Hirshleifer, D., & Welch, I. (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *American Economic Review*, 82(5), 992–1026.

- Brock, W. A., & Durlauf, S. N. (2002). A formal model of theory choice in science. In P. Mirowski & E.-M. Sent (Eds.), *Science bought and sold: Essays in the economics of science* (Chap. 11, pp. 341–361). University of Chicago Press.
- Clifford, P., & Sudbury, A. (1973). A model for spatial conflict. *Biometrika*, *60*(3), 581–588.
- DeMarzo, P. M., Vayanos, D., & Zwiebel, J. (2003). Persuasion bias, social influence, and uni-dimensional opinions. *Quarterly Journal of Economics*, *118*(3), 909–968.
- Durrett, R., & Steif, J. E. (1993). Fixation results for threshold voter systems. *Annals of Probability*, *21*(1), 232–247.
- Flocchini, P., Lodi, E., Luccio, F., Pagli, L., & Santoro, N. (1998). Irreversible dynamo in tori. In *Euro-Par98 parallel processing, Lecture Notes in Computer Science* (p. 1). Springer.
- Goldman, A. (1999). *Knowledge in a social world*. Oxford: Clarendon Press.
- Goles, E., & Olivos, J. (1980). Periodic behavior of generalized threshold functions. *Discrete Mathematics*, *30*, 187–189.
- Granville, A. (1991). On a paper of Agur, Fraenkel and Kleitman. *Discrete Mathematics*, *94*, 147–151.
- Hassin, Y., & Peleg, D. (2001). Distributed probabilistic polling and applications to proportionate agreement. *Information and Computation*, *171*, 248–268.
- Hegselmann, R., & Krause, U. (2002). Opinion dynamics and bounded confidence models, analysis, and simulation. *Journal of Artificial Societies and Social Simulation*, *5*(3).
- Hegselmann, R., & Krause, U. (2006). Truth and the cognitive division of labor: First steps toward computer aided social epistemology. *Journal of Artificial Societies and Social Simulation*, *9*(3).
- Holley, R. A., & Liggett, T. M. (1975). Ergodic theorems for weakly interacting infinite systems and the voter model. *The Annals of Probability*, *3*(4), 643–663.
- 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.
- Krčmar, R. (2001). On majority voting games in trees. In *IFSEM 2001: Theory and Practice of Informatics: 28th Conference on Current Trends in Theory and Practice of Informatics* (Piestany, Slovak Republic, November 24–December 1, 2001, p. 282).
- Latanž, B., & Nowak, A. (1997). Self-organizing social systems: Necessary and sufficient conditions for the emergence of clustering consolidation, and continuing diversification. *Progress in Communication Sciences*, *13*, 43–74.
- Moran, G. (1994a). Parametrization for stationary patterns of the r-majority operators on 0-1 sequences. *Discrete Mathematics*, *132*, 175–195.
- Moran, G. (1994b). The r-majority vote action on 0-1 sequences. *Discrete Mathematics*, *132*, 145–174.
- Moran, G. (1995). On the period-two-property of the majority operator in infinite graphs. *Transactions of the American Mathematical Society*, *347*(5), 1649–1667.
- Mustafa, N. H., & Pekev, A. (2004). Listen to your neighbors: How (not) to reach a consensus. *Journal of Discrete Mathematics*, *17*(4), 634–660.
- Nakata, T., Imahayashi, H., & Yamashita, M. (1999). Probabilistic local majority voting for the agreement problem on finite graphs. In E. A. T. Asano (Ed.), *COON '99 LNCS 1627* (pp. 330–338).
- Nakata, T., Imahayashi, H., & Yamashita, M. (2000). A probabilistic local majority polling game on weighted directed graphs with an application to the distributed agreement problem. *Networks*, *35*(4), 266–273.
- Peleg, D. (1998). Size bound for dynamic monopolies. *Discrete Applied Mathematics*, *86*, 263–273.
- Peleg, D. (2002). Local majorities, coalitions and monopolies in graphs: A review. *Theoretical Computer Science*, *282*, 231–257.
- Poljak, S., & Sura, M. (1983). On periodical behaviour in societies with symmetric influence. *Combinatorica*, *3*(1), 119–121.
- Sarkar, H. (2007). *Group Rationality in Scientific Research*. Cambridge: Cambridge University Press.
- Steif, J. E. (1994). The threshold voter automaton at a critical point. *Annals of Probability*, *22*(3), 1121–1139.
- Strevens, M. (2003). The role of the priority rule in scientific explanation. *Journal of Philosophy*, *100*(2), 55–79.
- Strevens, M. (2006). The role of the matthew effect in scientific explanation. *Social Studies in History and Philosophy of Science*, *37*, 159–170.
- Weisberg, M., & Muldoon, R. (forthcoming). Epistemic landscapes and the division of cognitive labor. *Philosophy of Science*.

-
- Welch, I. (1992). Sequential sales, learning, and cascades. *Journal of Finance*, 47(2), 695–732.
- Zollman, K. J. (2007a). The communication structure of epistemic communities. *Philosophy of Science*, 74(5), 574–587.
- Zollman, K. J. (2007b). The epistemic benefit of transient diversity. Manuscript.