Politics, Philosophy & Economics
http://ppe.sagepub.com/

Social network structure and the achievement of consensus

Kevin JS. Zollman

Politics Philosophy Economics 2012 11: 26 originally published online 18 October 2011
DOI: 10.1177/1470594X11416766

The online version of this article can be found at:
http://ppe.sagepub.com/content/11/1/26

Published by:
SAGE
http://www.sagepublications.com

On behalf of:
The Murphy Institute of Political Economy

Additional services and information for *Politics, Philosophy & Economics* can be found at:

Email Alerts: http://ppe.sagepub.com/cgi/alerts

Subscriptions: http://ppe.sagepub.com/subscriptions

Reprints: http://www.sagepub.com/journalsReprints.nav

Permissions: http://www.sagepub.com/journalsPermissions.nav

>> Version of Record - Feb 14, 2012

OnlineFirst Version of Record - Oct 18, 2011

What is This?

Downloaded from ppe.sagepub.com at CARNEGIE MELLON UNIV LIBRARY on May 10, 2012

Special Issue article on Complexity

Politics, Philosophy & Economics
11(1) 26–44
© The Author(s) 2012
Reprints and permissions:
sagepub.co.uk/journalsPermissions.nav
DOI: 10.1177/1470594X11416766
ppe.sagepub.com
SAGE

Social network structure and the achievement of consensus

Kevin JS. Zollman
Carnegie Mellon University, USA

Abstract
It is widely believed that bringing parties with differing opinions together to discuss their differences will help both in securing consensus and also in ensuring that this consensus closely approximates the truth. This paper investigates this presumption using two mathematical and computer simulation models. Ultimately, these models show that increased contact can be useful in securing both consensus and truth, but it is not always beneficial in this way. This suggests one should not, without qualification, support policies which increase interpersonal contact if one seeks to improve the epistemic performance of groups.

Keywords
Social epistemology, network epistemology, consensus, linear pooling, bounded confidence

One of the central problems in modern society is determining how best to integrate the wide variety of opinions that exist about almost any topic. Politically, it is critical to successfully integrate different opinions about what one should do regarding issues of redistribution and social justice. Pragmatically, organizations must take actions which appropriately represent a synthesis of the views of the members of such organizations. Scientifically, we want to find the best way for researchers to pool their opinions, a way that would most effectively lead us to true conclusions.

Consciously or unconsciously, individuals do (to some degree) integrate information from others. The process of 'social influence' has been extensively investigated by social psychologists, and its existence cannot be disputed. Solomon Asch, a famous early experimenter on social influence, regarded the extremity of social influence to be of serious concern. According to Asch (1955: 34):

Corresponding author:
Kevin JS. Zollman, Baker Hall 135, Pittsburgh PA, 15213-3890
Email: kzollman@andrew.cmu.edu

Downloaded from ppe.sagepub.com at CARNEGIE MELLON UNIV LIBRARY on May 10, 2012

> 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.

It is not so clear, however, that social influence is deleterious to our functioning as social beings. It allows for the formation of consensus, and can potentially represent one way that we can integrate variegated data about a particular situation quickly and easily. In an earlier article (Zollman, 2010b), I suggested that the process might serve to make communities of individuals more reliable. Utilizing a mathematical model of social influence I illustrated how communities of individuals who are informed (with noise) about some particular fact can be more reliable than the individuals themselves.

I also addressed another important question: given that social influence is, likely, an inevitable feature of human social interaction, what is the best form for that integration to take? Is it better for people to be in contact with more or fewer individuals?

This later questions is an important one. There are many political movements whose aim is to bring more people in contact with one another in order to increase dialogue about one issue or another. We are increasing the degree to which scientists connect with one another via the internet and through more journals and conferences. Organizations often attempt to foster interaction among their members. All of these policies have as a tacit assumption that a group is somehow made better by increasing interaction.

It would be a mistake to assume uncritically that increasing interaction would be always beneficial. Social processes such as belief change are undoubtedly complex systems. This is true both in the colloquial sense of that word (that is, complicated) and also in the technical sense (that is, exhibiting nonlinearity, sensitivity to initial conditions, and cyclic or chaotic behavior). Because social interaction in belief formation exhibits some of these properties of complex systems, we cannot hope to understand them from mere reflection alone. Nor can we generalize from a few empirical cases. Those features of complex systems which make them mathematically interesting thwart generalizing from a limited number of cases. For this reason, one must study these systems using mathematical or computer-simulation techniques.

Utilizing models of this type, my previous work has suggested that whether social interaction is productive depends on the underlying learning situation. When individuals are actively engaged in a process of testing different options (such as different technologies or different drug treatments) increased social interaction can, surprisingly, be very harmful (Bala and Goyal, 1998; Ellison and Fudenberg, 1995; Zollman, 2007, 2009, 2010a). This is because when individuals must seek out evidence there is some benefit to diversity of opinion, since it leads groups to seek out diverse bits of evidence.[1] Having highly connected social networks tends to retard this sort of diversity, and so has a distinct negative impact.

But when all the available information arrives before any social interaction takes place, increased interaction can be helpful (Zollman, 2010b). Groups tend to improve

Downloaded from ppe.sagepub.com at CARNEGIE MELLON UNIV LIBRARY on May 10, 2012

as the number of potential interactions increases. This later mathematical model was limited in two ways. First, it dealt only with beliefs that could be represented as one of two states: either I believe a proposition or not. Second, it considered only three potential structures for social interaction. This article attempts to remedy this shortcoming by offering another model which allows for consideration of a richer set of beliefs and more potential structures for social interaction. The models in this article better represent political opinions, probability judgments, or any judgment that can be represented by a point on a continuous spectrum.

But ultimately our question will be the same: what is the best form for social influence to take? Is it better to have more rather than less social influence? Are all policies that attempt to secure reliable consensus by increasing the amount of interaction within a group likely to be successful in their aims?

Depending on the situation, one might be interested in securing one of two goals. In many political situations one might be satisfied with consensus without regard for which particular opinion makes up that consensus. Here there is no external judge of 'correctness'. In other circumstances, however, we may not just want consensus but also a consensus which approximates the truth. In the discussion that follows, I will consider both measures of propriety.

This article provides two mathematical models for social influence, which will allow us to compare different types of social arrangements as to their ability to secure (correct) consensus. The first model, called 'linear pooling', is presented in Section 1. The second model is a generalization of the first and is presented in Section 2. The limitations of both of these models are discussed in Section 3 and Section 4 concludes.

1. Linear pooling

Perhaps the most straightforward model of belief change under social influence was devised by French (1956). French imagines that each individual has an opinion represented by a real number in some interval (here we will use [0,1]). Each individual is also influenced by another to a certain degree, again represented by some non-negative number. Individuals then update their beliefs by averaging their own belief with all those who have influence over them. All individuals do this simultaneously and thus arrive at new beliefs. This process is then repeated until eventually the forces of social influence are in equilibrium – no one further changes their beliefs.[2]

Individuals might average their belief by simply splitting the difference with those around them (more formally, by adopting the arithmetic mean of the beliefs of their friends). Alternatively, individuals might adopt a *weighted* average, where some individuals have a greater influence on their new belief than others.

1.1. Mere consensus

Under what conditions should we expect a community of individuals of this sort to reach a consensus? Suppose that we represent influence by using a social network. Each individual is represented by a dot (or vertex) and if individual i is influenced by individual j to any degree greater than zero, we draw an arrow beginning at i and ending at j. Once we

Downloaded from ppe.sagepub.com at CARNEGIE MELLON UNIV LIBRARY on May 10, 2012

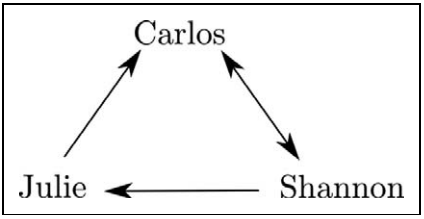

Figure 1. An example of an influence graph which converges in French's model to a single consensus opinion

have done this for every individual we now have a directed graph. If this graph is *connected* (that is, we can begin at any individual and, following the directions of the arrows, get to any other individual in some number of steps) and if at least one person is influenced by herself, the group as a whole will, in the limit, converge to a single consensus opinion (DeGroot, 1974). (While sufficient, this is not a necessary condition. Some other network structures will converge, and others will converge if the initial beliefs have the appropriate structure (see Berger, 1981).)

Consider a simple example. Suppose a community of three individuals: Carlos, Julie, and Shannon. Carlos is influenced by Shannon, but not by Julie. Julie is influenced by Carlos, and Shannon is influenced by both Julie and Carlos. This group is pictured in Figure 1. One can see that one can get from any individual to any other by following the arrows. As a result, this group will converge to a consensus. The final consensus is determined by the degree to which each individual is influenced by the others, but can be calculated by a relatively simple mathematical procedure. The details of this procedure can be found in the Appendix.

If our interest is the creation of consensus, what does this tell us? If we have two groups that appear to be at odds, we would like to bring them into contact. It is sufficient to have a single person from one group who is influenced by one person in the other group, but more cannot hurt.[3] In this way, French's model underwrites those who believe that our political institutions would be improved by increased communication.

1.2. Correct consensus

For some decisions, consensus may be all we are interested in. For instance, in some democratic decision making, the 'correctness' of a decision may merely depend on it being the option agreed on by everyone (or at least the majority).[4] In such cases, all we desire is to construct a situation in which consensus is achieved.

However, in many cases there is an independent truth to be tracked. We want to create a situation not only in which consensus is achieved, but also in which the consensus is as close as possible to the truth. For this circumstance we will need to add some details to French's model in order to make the analysis more tractable. First, we will restrict our notion of social influence. Instead of allowing arbitrary degrees of influence, we will impose some constraints. Suppose every individual has a group of friends (we will

Downloaded from ppe.sagepub.com at CARNEGIE MELLON UNIV LIBRARY on May 10, 2012

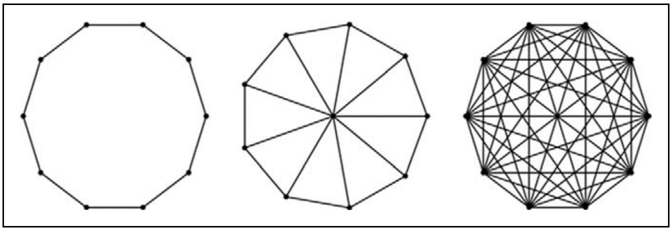

Figure 2. Three graphs: the cycle, the wheel, and the complete graph. The left-hand and right-hand graphs are regular graphs, while the center graph is not

assume friendship is symmetric, that is, if I am your friend, then you are my friend). Suppose, too, that each individual is influenced to the same degree by all of her friends.[5] This allows us to represent social influence as a undirected graph, such as those pictured in Figure 2.

Now suppose that each individual is given some initial 'signal' about the true value of some parameter. Perhaps every individual performs an independent experiment to estimate the mean of some distribution or perhaps they are each given different information about the real value of some underlying probability. Whatever the situation being modeled, we will presume that the expected value for each individual is the truth.

Given this model, we can now ask what structure of social influence produces a consensus that is closest to the truth. If each individual is equally reliable, the structure of social influence should be *regular*, that is, each individual should have exactly the same number of friends as everyone else (proofs can be found in the Appendix, and see also DeMarzo et al. (2003) and Golub and Jackson (2010)).[6] This ensures that no individual has undue influence over the final consensus estimate. Consider the three graphs in Figure 2. Here the left-hand and right-hand graphs are both regular. The left-hand graph is the cycle, in which every individual has exactly two friends (plus herself). The right-hand graph is the complete graph, in which everyone has everyone else as a friend. Surprisingly, perhaps, judged from the ultimate consensus properties both are equally good. These are in contrast to the middle graph from Figure 2, the wheel, in which a single individual is friends with everyone. Here, the person in the center exerts disproportionate influence on the final consensus and so unduly skews the result in favor of her own opinion.

There is an important difference between the cycle and the complete graph. The complete graph is much faster at reaching consensus – it does so in one step. The cycle, on the other hand, takes literally forever to reach consensus, although it gets closer on every step. So if one is concerned not just with ultimate consensus, but also with speed, then one might prefer the complete graph (Harray, 1959).

These results make important qualifications to the conclusions of Section 1.1. There it was the case that the addition of any new friends could not hurt in securing consensus, but here it clearly can hurt. If one begins with a group that is arranged in a cycle and adds a single link by introducing two new people, one has now made the group less reliable. This suggests that if we have an idea about truth toward which we want the group to aim, then bringing individuals into contact with one another should not be regarded as always productive.

Downloaded from ppe.sagepub.com at CARNEGIE MELLON UNIV LIBRARY on May 10, 2012

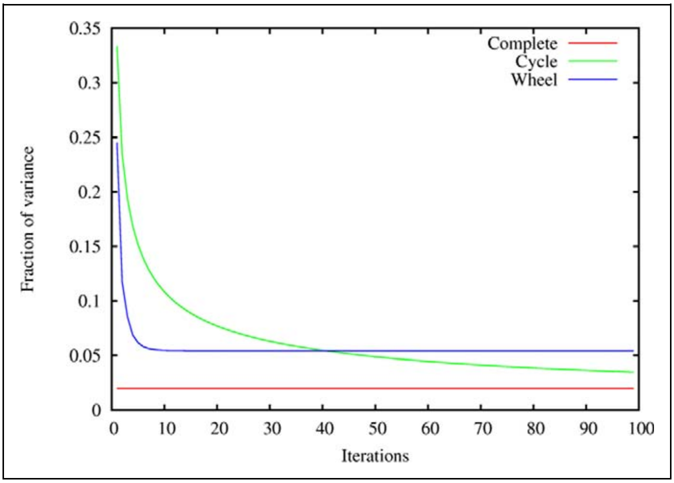

Figure 3. The short-run behavior of the cycle, the wheel, and the complete graph using linear pooling

A comparison of the short- and long-run behavior of these two graphs can be found in Figure 3. Here the x-axis represents the number of iterations of the averaging procedure and the y-axis represents how far (on average) the final judgment is from the truth.[7] Looking only at very short-term behavior (on the far the left of the graph), the addition of edges helps even if it reduces the regularity of the graph. But after a sufficient amount of time, the unequal number of connections in the wheel results in a more skewed estimate of the truth than does the cycle, which approaches the reliability of the complete graph.

What if individuals are unequally reliable? Perhaps one person is given two signals or a larger sample. In which case, we want the friendship properties to represent this difference in the quality of information. If one person is twice as reliable as everyone else, we want that person to have twice as many friends (counting herself). Three examples are pictured in Figure 4. In each of these examples, the addition of *any number of edges whatsoever* would reduce the reliability of the group. Again, we might not want to introduce two new people here.

One might think that requiring all friends be treated equally is too stringent, since people are influenced to different degrees by their different friends. Space prevents a detailed discussion of this situation here, but interested readers should consult Golub and Jackson (2010) and Hartmann et al. (2009), who consider exactly this point.

2. Bounded confidence

Some authors have complained that French's model is an inadequate empirical model because it fails to explain the persistence of disagreement in situations of social influence (Abelson, 1964; Friedkin and Johnsen, 1999). This led to the development of a similar model in which individuals' social influence depends on the distance of their beliefs

Downloaded from ppe.sagepub.com at CARNEGIE MELLON UNIV LIBRARY on May 10, 2012

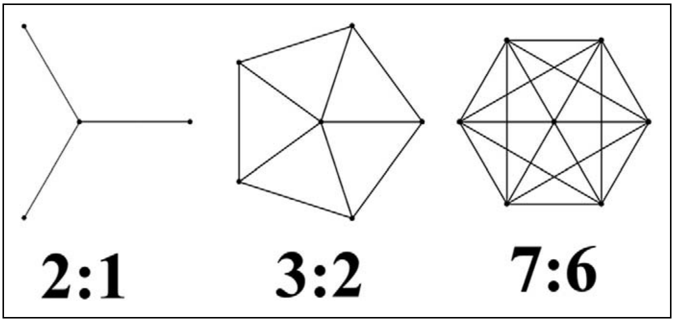

Figure 4. Maximally connected optimal graphs for heterogeneous reliability, in which the proportions represent the number of samples by a very reliable individual compared to the number of samples by the remaining individuals

from one another. The bounded confidence model retains the central thrust of French's model while allowing for disagreement even in connected graphs (Hegselmann and Krause, 2002; Krause, 2000).[8]

Like French's model individuals are endowed with initial beliefs and are influenced by others. Along the lines of our simplification above, it is assumed that individuals have a set of friends who all have the potential to be equally influential. However, each individual whittles down her set of friends on the basis of similarity of opinion. A friend whose opinion differs by too great a degree is ignored. Of those friends who are sufficiently close in opinion, the agent averages her own belief with theirs, just as in French's model. The largest difference in opinion a person is willing to tolerate is called the 'tolerance parameter'. We will place individuals in a social network and assign them all the same tolerance parameter.

2.1. Mere consensus

Graph structure will also influence the degree to which consensus is achieved in this model. Like French's model, in unconnected graphs we cannot be assured that consensus will be achieved. Unlike French's model, even connected graphs might retain persistent disagreement. Consider, for example, a community composed of two individuals, each of whom considers those within 0.5 of their own belief credible and thus worthy of consideration. If one individual has a belief of 0.1 and the other 0.9, the group will never reach a consensus. On the other hand, if one has 0.4 and the other 0.6, they will.

It seems likely that the connectedness of the graph will have an influence on the degree of consensus which is achieved. Hegselmann and Krause (2002) show that in a complete graph the tolerance parameter puts an upper bound on the number of distinct groups which persist in the limit. In particular, it is impossible for there to be more than $1/\delta$ distinct groups in the long run (where δ is the tolerance parameter). However, in the cycle graph (also a regular graph) there can be as many distinct groups as individuals.[9] In terms of worst-case scenarios, the complete graph appears far better than the cycle.

Downloaded from ppe.sagepub.com at CARNEGIE MELLON UNIV LIBRARY on May 10, 2012

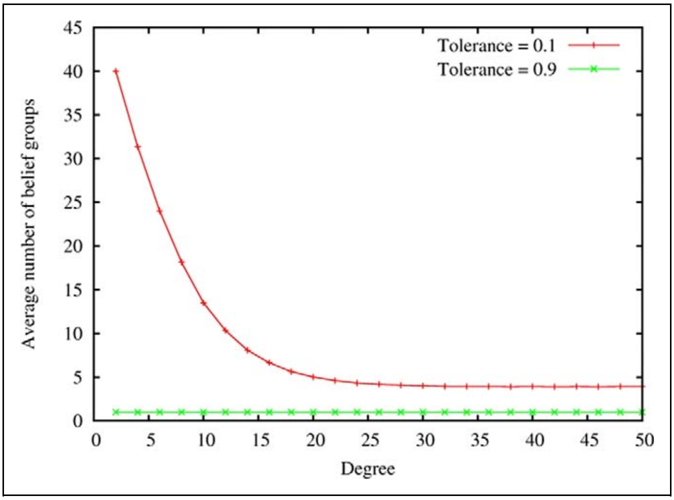

Figure 5. Simulation results illustrating the average number of belief groups for various *x*-cycles and two extreme degrees of tolerance

But is the complete graph superior when we compare situations other than the worst case? Here analytical solutions will be too difficult; we must turn to simulation. We will first consider a family of graphs known as the *x*-cycles. Individuals are arranged in a circle and are connected to their *x* nearest neighbors ($x/2$ on the right and $x/2$ on the left). The cycle pictured on the left in Figure 2 is a 2-cycle. The 4-cycle involves connecting everyone to their neighbors' neighbors, while the 6-cycle begins with the 2-cycle and connects everyone to their neighbors' neighbors and their neighbors' neighbors' neighbors, and so on. This family of graphs is useful because every member is regular (everyone has the same number of neighbors) and one can move from a graph with very few connections to the complete graph.

Figure 5 shows simulation results that illustrate the relationship between connectivity and the number of distinct opinion groups. On the x-axis is the degree of *x*-cycle which is being considered and on the y-axis is the average number of distinct opinion groups which remained after 1000 iterations of averaging.[10] If the number of distinct opinion groups is one, consensus is always achieved. Higher numbers indicate either that consensus is achieved less often or that it is never achieved at all.[11] For low values of the tolerance parameter, connectedness helps to reduce the number of distinct opinion groups – this brings us closer to consensus in this sense. But the benefit of connectedness appears to decrease rapidly. Once one is connected to half the network (degree 24 or 26) very little is gained from being connected to more. For high tolerance values, however, connectedness has little effect. This should not be surprising since when the tolerance level is 1.0, the bounded confidence and linear pooling models are equivalent. The last, also unsurprising, result is that higher tolerance levels (that is, having groups of open-minded individuals) tend to assist in the achievement of consensus in all the graphs studied here.

What about regularity? In the linear pooling model, regularity had no effect whatsoever on the ultimate achievement of consensus. Here it appears to make a difference for

Downloaded from ppe.sagepub.com at CARNEGIE MELLON UNIV LIBRARY on May 10, 2012

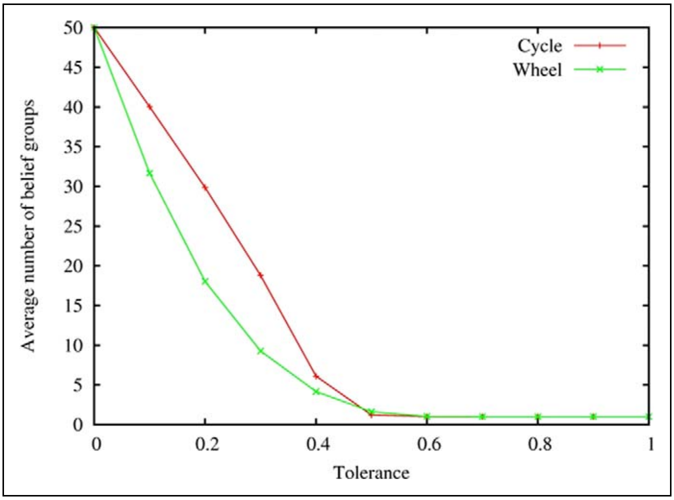

Figure 6. Simulation results illustrating the average number of belief groups for the cycle and the wheel and various tolerance levels

relatively small tolerance levels. Figure 6 compares the wheel and the cycle for various tolerance levels. For groups of relatively close-minded individuals, increasing the number of connections helps to secure consensus even when that connectivity is achieved by introducing irregularity (that is, the wheel performs better than the cycle). This makes intuitive sense, since the more individuals that one interacts with, the more likely it is that one will interact with someone who acts as a bridge between oneself and another. This is true even if the other person is very central.

The presence of central individuals might create consensus by influencing everyone else to take their positions. While central individuals might help to create consensus, they might do so by making the group as a whole less reliable (when there is a well-defined sense of reliability). We turn to this possibility in the next section.

2.2. Correct consensus

The previous section focused only on the achievement of consensus, where this was measured by counting the number of distinct belief groups which existed after a sufficiently long simulation. As before, we might also be interested in measuring the distance from the truth for certain groups. Using the bounded confidence model, what is the effect of connectivity and regularity on the truth-seeking properties of the groups?

Interestingly, connectivity is most beneficial for those who are somewhat open-minded, but not totally open-minded. Figure 7 illustrates the results from the x-cycles for three different tolerance levels. The y-axis represents average squared distance from the truth, and as a result, lower scores represent groups who, on average, more closely approximate the truth. The top and bottom line represent groups made up of individuals with very small or very large tolerance parameters. Although the former is helped by increased connectivity, the benefit is far less than the effect connectivity has on those

Downloaded from ppe.sagepub.com at CARNEGIE MELLON UNIV LIBRARY on May 10, 2012

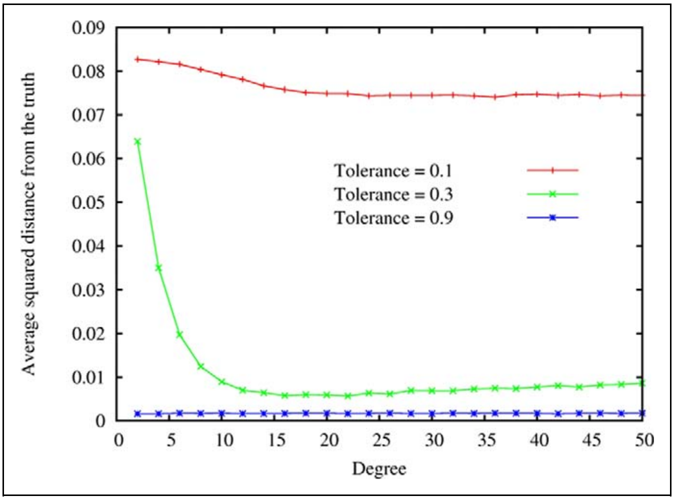

Figure 7. Simulation results illustrating the average squared distance from the truth for various cycles and three tolerance levels

with a tolerance parameter of 0.3 (the middle line). The moral of this result is that certain types of groups are likely to benefit from increased contact, namely, those groups comprising individuals who are neither extremely open- or extremely close-minded. Those that are too close-minded are unlikely to listen to the new people that they meet and those that are too open-minded will eventually receive the information via another source anyway.

In the linear pooling model, regularity was important because it prevented a single person from dominating the consensus – it gave each person equal say. Again, we should expect this to be true for high tolerance levels, and this is what we find through simulation (as illustrated in Figure 8). For high values of the tolerance parameter (above 0.4), we see that groups arranged in the cycle do better than groups arranged in the wheel. The conclusion here is the same as for the linear pooling model: increasing connections arbitrarily might have a negative impact on the reliability of the group.

Surprisingly, however, this fact does not hold for low values of the tolerance parameter. When individuals are relatively close-minded, it is better to have a central individual, even though that person's opinion might have an overly high influence on the final consensus. This is achieved because the central individual can act as a bridge between two different people. Jake and Carlos might have opinions outside of the other's tolerance parameter, and so might not integrate the other's information. But if Julie occupies some middle ground between Jake and Carlos (but within each of their tolerance levels), she might bring them close enough so that they do end up integrating information from one another.[12]

3. Limitations of these models

Underlying these models are a series of assumptions, many of which are unrealistic. One might, therefore, be interested in knowing the degree to which these assumptions are

Downloaded from ppe.sagepub.com at CARNEGIE MELLON UNIV LIBRARY on May 10, 2012

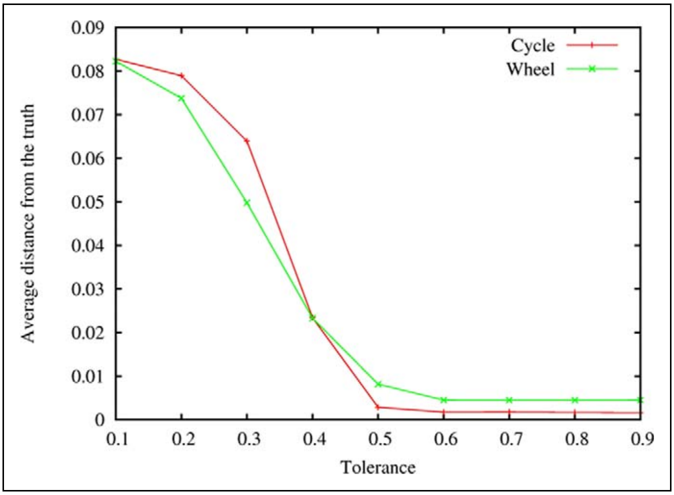

Figure 8. Simulation results illustrating the average distance from the truth for the cycle and the wheel and various tolerance levels

important. First among them is the very simple model of social influence represented by repeated averaging. Bringing individuals together to discuss various issues is not supposed to get them simply to 'average' their opinions with one another, but is, instead, supposed to be an open discussion of the issues in which people share the reasons for holding their beliefs. No one would deny that, at least sometimes, social interaction involves a much more complex method by which individual beliefs are changed. However, this model is not supposed to be about those details, but is, instead, a macroscopic model of individual belief change. Those authors who use these models hope that it captures, as a first approximation, how beliefs are changed, and therefore the conclusions hold in general for groups that change their beliefs in this more complicated way.[13]

Second, we have assumed that each individual is equally influenced by all her friends. This is unlikely to be true in the real world, but it represents a starting point for analysis. Other authors have considered departures from this assumption (see Demarzo et al., 2003; Golub and Jackson, 2010; Hartmann et al., 2009).

A third assumption is that the network is static, that is, one communicates with the same individuals over time. One might imagine modeling the process by which friendships change over time. How they might change will depend on the circumstance and the people involved. Holme and Newman (2006) consider a situation in which individuals tend to connect with those who have similar opinions. This is certainly a reasonable model, and has some affinity with the assumptions that underlie the bounded confidence model. But, there may be circumstances in which people seek out those with different opinions from their own.

In this vein, one might be interested to know what networks might arise when individuals are allowed to form their own networks. For instance, if individuals were rewarded for the reliability of their final consensus opinion and each could chose his group of friends, what network might form? If adding additional friends were free, in all

Downloaded from ppe.sagepub.com at CARNEGIE MELLON UNIV LIBRARY on May 10, 2012

situations the complete graph would seem a likely end point. But if there was a cost to adding friends, would this still be the case? In some cases, the cycle might be the best end point, but it might be difficult to achieve (it would require significant coordination among the players). Although models of network formation and learning have been considered (see Goyal, 2007; Jackson, 2008), none have addressed learning problems of exactly this sort.

A final assumption made by this model is that we are dealing with a situation in which all the information has already arrived – no new information arrives during the process of deliberation. Situations of this sort are not uncommon; risk assessment, jury deliberation, and departmental hiring decisions all have this form. But, of course, many situations of deliberation are not – new information can arise which might influence some people and not others.

There are various ways that this might be incorporated into a model of this form. First, one might simply imagine that there is some constant ‘pull’ toward the truth. In each round, some (or all) individuals average with their neighbors and are also pulled in the direction of the truth. This model has been investigated (see Douven and Riegler, 2010; Hegselmann and Krause, 2006; Riegler and Douven, 2009) and has some interesting properties.

This way of modeling the arrival of new information also has an important idealizing assumption: the information which arrives is independent of the beliefs of the inquirers. Again, this is not always the case. Sometimes the type of information we solicit (through experiments, research, and so on) depends on what we think. In previous research (Zollman, 2007, 2009, 2010a), I have considered a model of this form, one in which individuals need to seek out evidence. In these cases, connectivity is a bad thing (whether regular or not).

Comparing the result here with this earlier work is insightful. It illustrates that even some of the most basic questions about how to design social institutions (such as how much people should communicate) can depend radically on the underlying structure of the problem such institutions are trying to solve. In the case of the model presented here, in which individuals are trying to integrate already gathered information, highly connected and regular networks appear to be best, all things considered, while in other models highly connected graphs are the worst (Zollman, 2010a).

Although the models considered here are unrealistic in various ways, modifying the model to make it as realistic as is feasible may not be productive. Models which are as complex as the phenomena they study also have the property of being as impenetrable as the real phenomena, making the models useless as heuristic devices for understanding the social world. Instead, the step-by-step process whereby models are made slightly more realistic (and then analyzed carefully before proceeding) provides, in this author’s mind, the most effective way to proceed in understanding complex social phenomena such as the formation of a consensus.

4. Conclusion

Given that individuals are subject to social influence, how should they interact so as to form groups which achieve consensus and are reliable? We seem to have two features of

Downloaded from ppe.sagepub.com at CARNEGIE MELLON UNIV LIBRARY on May 10, 2012

social interaction which tend to make groups better in the sort of situation considered here: connectivity and regularity. While not inconsistent, these two features do not always select the same thing. One can take a graph and increase its connectivity, but reduce its regularity and vice versa.

So, when in conflict, which virtues are dominant? Here we have different conclusions for different circumstances. When individuals are relatively close-minded, one should try to bring individuals together regardless of whether doing so introduces irregularity. Bringing new people together moves the group as a whole closer to a consensus and also makes the group more reliable.

If, on the other hand, the group is made up of relatively open-minded individuals, the answer is more complex. If consensus is all that is desired, without concern for correctness, then again connectivity is the sole good. It will better approximate consensus, and even when a consensus is guaranteed, it reduces the time needed to reach that consensus. However, if one is not just concerned with the achievement of consensus, but also the correctness of the final result, one must balance connectedness with regularity. Simply putting new people in contact with one another will not always result in better outcomes. This suggests that in at least some situations, one ought to exercise caution when attempting to increase communication – it might not always serve to make the group as a whole more reliable.

Beyond these individual recommendations for improving the reliability of social groups, the conclusions of this article show how dangerous limited analysis could be. From a superficial analysis, one might conclude that connectivity is always good and seek to bring new people into contact with one another. In some circumstances, this could create epistemically worse groups if one cannot approach a complete graph. This sort of 'nonlinearity' should not strike those familiar with complex systems as surprising, but nonetheless it is an important consequence to keep in mind when considering the epistemic properties of groups.

Appendix

Each individual is endowed with an initial belief .[14] For each pair of individuals i and j there is a weight w_{ij} which represents the influence that j has over i. (For simplicity we will define i's influence on herself as well, but this can just be thought of as i's resistance to others' influence.[15]) So that our beliefs remain contained in [0,1], we will restrict our weights to the same interval, and also require that for each i, $\Sigma_j\, w_{ij} = 1$. Suppose that individuals simultaneously update their beliefs by weighted averaging. Their new belief is given by

$$b_i^1 = \sum_j w_{ij} b_j^0 \qquad (1)$$

This single iteration will not always result in a stable belief. Since i has moved in response to j's old belief, but j has also changed in response to others, the new beliefs can be further subjected to social influence. Thus, we can define an iterative process, in which later b_i^n are constructed by the same method of synchronous averaging.

Downloaded from ppe.sagepub.com at CARNEGIE MELLON UNIV LIBRARY on May 10, 2012

Consider the matrix W constructed by the weights, w_{ij}, and the matrix B constructed by the beliefs, b_i^0. One step of this process is simply the matrix multiplication of these two matrices, WB. This gives us the matrix of beliefs represented by Equation 1. A second step of this process is just a successive multiplication by W, W(WB). So, the *n*th step is just $W^n B$. This has some nice formal properties which make the analysis of this model relatively easy. If there is a matrix L such that $L = \lim_{n \to \infty} W^n$, then this iterative process converges to LB.

Because of the conditions placed on the weights, W is a Markov chain. Thus, many of the results regarding the existence of stationary distributions in Markov chains can be applied here as well.[16] If W is ergodic we can be guaranteed the existence of a stationary distribution. It is sufficient that the influence graph (discussed above) be connected and at least one person is influenced to some positive degree by herself.

We will now define two ways of deriving the weight matrix from an underlying social network. Let $G = \langle N, E \rangle$ be a graph, composed of a set of nodes, *N*, and a set of edges, *E*. Let $\mathcal{N}_i$ represent *i*'s neighborhood (including *i*). We can now define the weight matrix, W, this way,

$$w_{ij} = \begin{cases} 1/|\mathcal{N}_i| & \text{if} j \in \mathcal{N}_i \\ 0 & \text{Otherwise} \end{cases} \tag{2}$$

This creates a perfectly symmetric situation based on the underlying graph. Each individual is equally influenced by himself and his neighbors and is not directly influence by those to whom he is not connected.

The stationary distribution of a Markov chain can be computed by solving the following set of equations:

$$\pi_i = \sum_j \pi_j w_{ji}$$

$$1 = \sum_j \pi_i$$

The vector $\Pi = \langle \pi_i \rangle$ represents the stationary distribution of the Markov chain W, and thus represents the weight each individual is given in determining the final consensus of the group. Given that we have begun with an underlying network, we are able to show a relationship between the stationary distribution and the underlying network structure from which we derived the weight matrix. This result is expressed in the following proposition (DeMarzo et al., 2003).

Proposition 1: For any network assigning weights according to Equation 2, the stationary distribution Π of W is $\pi_i = \frac{|\mathcal{N}_i|}{2|E| + n}$, where $|E|$ is the number of edges in the graph and *n* is the number of nodes.

Proof: Recall that each entry in W is either $\frac{1}{\mathcal{N}_i}$ or 0. So,

Downloaded from ppe.sagepub.com at CARNEGIE MELLON UNIV LIBRARY on May 10, 2012

$$\pi_i = \sum_{j \in \mathcal{N}_i} \pi_j \frac{1}{|\mathcal{N}_j|}$$

Since $i \in \mathcal{N}_i$, we can subtract this term from both sides, yielding

$$\frac{|\mathcal{N}_i| - 1}{|\mathcal{N}_i|} \pi_i = \sum_{j \in \mathcal{N}_i / i} \pi_j \frac{1}{|\mathcal{N}_j|} \tag{3}$$

Suppose that $\pi_j = \frac{|\mathcal{N}_j|}{2|E| + n}$, to show that this satisfies this constraint of a stationary distribution. Substituting this value for π_j in Equation 3 and multiplying by the reciprocal yields:

$$\pi_i = \frac{|\mathcal{N}_i|}{|\mathcal{N}_i| - 1} \sum_{j \in \mathcal{N}_i / i} \left(\frac{|\mathcal{N}_j|}{2|E| + n} \right) \left(\frac{1}{|\mathcal{N}_j|} \right)$$

This reduces to:

$$\begin{aligned} \pi_i &= \frac{|\mathcal{N}_j|}{|\mathcal{N}_j| - 1} \sum_{j \in \mathcal{N}_i / i} \frac{1}{2|E| + n} \\ &= \frac{|\mathcal{N}_j|}{2|E| + n} \end{aligned}$$

Since $\sum_i |\mathcal{N}_i| = 2|E| + n$, this satisfies the last constraint that the sum of all π_i's equals 1.

One might be concerned that the individuals in the system represented by Equation 2 are overly deferential to others' beliefs. Instead, we might consider a more asymmetric method of assigning weights. Suppose there is a single number $s \in (0, 1)$, which every individual assigns to herself and then divides the remaining weights evenly among her remaining neighbors. This gives us the following weight assignments:

$$w_{ij} = \begin{cases} s & \text{If} i = j \\ (1 - s)/(|\mathcal{N}_i| - 1) & \text{If} j \in \mathcal{N}_i / i \\ 0 & \text{Otherwise} \end{cases} \tag{4}$$

We can prove a similar result for this new weight assignment.

Proposition 2: For any network assigning weights according to Equation 4, the stationary distribution Π of W is $\pi_i = \frac{|\mathcal{N}_i| - 1}{2|E|}$.

Proof. For every individual i,

Downloaded from ppe.sagepub.com at CARNEGIE MELLON UNIV LIBRARY on May 10, 2012

$$\pi_i = s\pi_i + \sum_{j \in \mathcal{N}_i/i} \pi_j \frac{1-s}{|\mathcal{N}_j| - 1}$$

which reduces to

$$(1-s)\pi_i = \sum_{j \in \mathcal{N}_i/i} \pi_j \frac{1-s}{|\mathcal{N}_j| - 1} \tag{5}$$

Again, we will show that $\pi_i = \frac{|\mathcal{N}_j|-1}{2|E|}$ satisfies the constraints by substituting it into Equation 5. This substitution results in

$$\begin{aligned}(1-s)\pi_i &= \sum_{j \in \mathcal{N}_i/i} \left(\frac{|\mathcal{N}_j|-1}{2|E|}\right)\left(\frac{1-s}{|\mathcal{N}_j|-1}\right) \\ \pi_i &= \frac{1}{(1-s)} \sum_{j \in \mathcal{N}_j/i} \left(\frac{|\mathcal{N}_j|-1}{2|E|}\right)\left(\frac{1-s}{|\mathcal{N}_j|-1}\right) \\ &= \sum_{j \in \mathcal{N}_j/i} \frac{1}{2|E|} \\ &= \frac{|\mathcal{N}_i|-1}{2|E|}\end{aligned}$$

This also satisfies the constraint that $\Sigma_j \pi_i = 1$, and so, as a result, represents the stationary distribution of W.

Propositions 1 and 2 show that the final weight given to each individual depends on the size of her neighborhood and on the number of edges in the graph. Given a fixed graph, more highly connected individuals will receive higher final weights as compared to those who are less connected. Propositions 1 and 2 entail that all regular graphs will have identical stationary distributions, since in each case each node will be given $1/n$ weight.[17] As a result, they will have identical limiting behavior.

These propositions also establish a relatively simple correctness result. If one has n independent and identically distributed samples from a distribution, one minimizes the squared distance from the mean just in case one averages the samples with equal weights. So if each individual's initial belief is drawn from some distribution, the network which does the best is one where the consensus weights assign $1/n$ weight to each individual. As a result, the regular networks are best when individuals are all equally reliable.

If we model differences in reliability as differences in the number of initial samples, the edges in the graph should be distributed so that each individual draw is given equal weight. This generates the graphs pictured in Figure 4.

Notes

The author would like to thank Brian Skyrms, Jeff Barrett, Rainer Hegselmann, John Rapalino, Kai Spiekermann, Kyle Stanford, Elliott Wagner, Andrew Williams, and the participants of the PPE Conference on Complexity for their useful comments and discussion.

1. In these models, I assume that everyone is equally capable of finding the right answer. Diversity might have other benefits as well; for instance, it might help in groups in which individuals differ in terms of the problems they can solve (Hong and Page, 2001, 2004).

Downloaded from ppe.sagepub.com at CARNEGIE MELLON UNIV LIBRARY on May 10, 2012

2. This model has been utilized as a normative model for incorporating the beliefs of others into one's beliefs by DeGroot (1974) and Lehrer and Wagner (1981). I will leave this issue about the normative correctness of this model to one side and instead focus on this model as an idealization of the process of social influence. For those who regard this model as a normative model for belief change, this article will have additional relevance.
3. The interested reader should consult Grim et al. (2010) for a discussion of the benefits of increasing connections between disparate groups when one is considering not just whether a group converges, but the speed at which it converges as well.
4. I do not wish to contradict the view of 'epistemic democrats', who believe that democracy is valuable because of its ability to reach true conclusions about facts and issues of justice. But they would certainly agree that at least some decisions (such as deciding on a dinner location for a group) are 'correct' only insofar as they are the consensus decision.
5. We can allow an individual to be influenced by herself to a greater (or smaller) degree than by her friends. The details of the model are provided in the Appendix.
6. This result does not depend on the distribution from which the initial signal is drawn. For any distribution with a finite variance, so long as each individual receives a single independent draw from that same distribution, the best network is regular.
7. Distance from the truth here is measured in fractions of variance of the distribution from which the initial beliefs are drawn. So, 0.1 means that the distribution of the consensus estimates will have one-tenth the variance of the distribution from which the initial beliefs are drawn. This figure is an analytical result, not the result of simulation, and so applies to any distribution with finite variance.
8. There are other ways that one might account for this intuition. One might allow individuals to assign less weight, but not zero weight, to those who are further away. See Jackson (2008: Ch. 8) for a discussion.
9. This applies to graphs with an even number of individuals. For graphs with an odd number of individuals and tolerance parameters greater than 0.5, there can only be n–1 distinct groups (where n is the number of individuals in the graph).
10. In all the simulations presented here, individual's initial beliefs are drawn from a uniform distribution over [0,1]. All groups have 50 individuals and the simulation results presented are for 5000 different simulations each of 1000 iterations of the averaging procedure. Simulations were written and run using the NetLogo modeling software. The simulations are available from the author's website at http://www.andrew.cmu.edu/user/kzollman.
11. In the case in which consensus is never achieved, it still seems better to have fewer rather than more belief groups. After all, resolving two competing opinions is often easer than resolving ten.
12. All of the results reported here are for simulations in which the initial belief was drawn from a uniform distribution over [0,1] (where the truth was 0.5). The qualitative results about which networks and tolerance parameters fare better appears robust to modifying the truth parameter. Simulations were checked when the initial beliefs were drawn at random from a beta distribution with parameters 2 and 20, 2 and 2, and 20 and 2 (where the truth is 0.1, 0.5, and 0.9, respectively), and the qualitative relationships remained the same.
13. See Hegselmann and Krause (2006) for a discussion of this issue.
14. For our purposes we will restrict ourselves to beliefs represented as points in a one-dimensional continuous space. DeGroot (1974) suggests that extending this model to

Downloaded from ppe.sagepub.com at CARNEGIE MELLON UNIV LIBRARY on May 10, 2012

probability density functions would be trivial, although in the case of conjugate distributions, such as the Beta distribution, there are some difficult choices to be made regarding what exactly is averaged by the agents (Winkler, 1968).

15. Here, we suppose that *i*'s influence over *j* remains constant over time. A model in which this changes over time has been analyzed in a slightly different context by Chatterjee and Seneta (1977), Lorenz (2005), and Cohen et al. (1986). Because we are assuming a fixed social network, we will stick with French's initial assumption.
16. A stationary distribution in a Markov chain is just L (as defined above). Depending on the weights, L need not always exist. The relationship between this model and Markov chains was first noticed by Harray (1959), who applied some theorems about Markov chains to French's model in order to determine when consensus would occur. This relationship represents a central part of DeGroot's discussion of this model (1974).
17. This fact was proved for the case of uniform weighting by Harray (1959).

References

Abelson RP (1964) Mathematical models of the distribution of attitudes under controversy. In: Frederiksen N, Gulliksen H (eds) *Contributions to Mathematical Psychology*. New York: Holt, Rinehart, and Winston.

Asch SE (1955) Opinions and social pressure. *Scientific American* 193(5): 31–5.

Bala V and Goyal S (1998) Learning from neighbours. *Review of Economic Studies* 65: 565–621.

Berger RL (1981) A necessary and sufficient condition for reaching a consensus using DeGroot's method. *Journal of the American Statistical Association* 76(374): 415–8.

Chatterjee S and Seneta E (1977) Toward consensus: Some convergence theorems on repeated averaging. *Journal of Applied Probability* 14: 89–97.

Cohen JE, Hajnal J and Newman CM (1986) Approaching consensus can be delicate when positions harden. *Stochastic Processes and their Applications* 22: 315–22.

DeGroot MH (1974) Reaching a consensus. *Journal of the American Statistical Association* 69(345): 118–21.

DeMarzo PM, Vayanos D and Zwiebel J (2003) Persuasion bias, social influence, and unidimensional opinions. *Quarterly Journal of Economics* 118(3): 909–68.

Douven I and Riegler A (2010) Extending the Hegselmann-Krause model I. *Logic Journal of the IGPL* 18(2): 323–35.

Ellison G and Fudenberg D (1995) Word-of-mouth communication and social learning. *Quarterly Journal of Economics* 110(1): 93–125.

French JRJ (1956) A formal theory of social power. *Psychological Review* 63(3): 181–94.

Friedkin NE and Johnsen E (1999) Social influence networks and opinion change. *Advances in Group Processes* 16: 1–29.

Golub B and Jackson MO (2010) Naive learning in social networks and the wisdom of crowds. *American Economic Journal: Microeconomics* 2(1): 112–49.

Goyal S (2007) *Connections: An Introduction to the Economics of Networks*. Princeton, NJ: Princeton University Press.

Grim P, Reade C, Singer DJ, et al. (2010) What you believe travels differently: Network homophily and the contrast between information and infection dynamics. *Group for Logic and Formal Semantics Research Report* (10-03).

Downloaded from ppe.sagepub.com at CARNEGIE MELLON UNIV LIBRARY on May 10, 2012

Harray F (1959) A criterion for unanimity in French's theory of social power. In: Cartwright D (ed.) *Studies in Social Power*. Ann Arbor: University of Michigan Press.

Hartmann S, Martini C and Sprenger J (2009) Consensual decision-making among epistemic peers. *Episteme* 6(2): 110–29.

Hegselmann R and Krause U (2002) Opinion dynamics and bounded confidence models, analysis, and simulation. *Journal of Artificial Societies and Social Simulation* 5(3).

Hegselmann R and Krause U (2006) Truth and the cognitive division of labour: First steps toward a computer aided social epistemology. *Journal of Artificial Societies and Social Simulation* 9(3).

Holme P and Newman MEJ (2006) Nonequilibrium phase transition in the coevolution of networks and opinions. *Physical Review E* 74(056108): 1–5.

Hong L and Page SE (2001) Problem solving by heterogeneous agents. *Journal of Economic Theory* 97: 123–63.

Hong L and Page SE (2004) Groups of diverse problem solvers can outperform groups of high-ability problem solvers. *Proceedings of the National Academy of Sciences* 101(46): 16,385–9.

Jackson MO (2008) *Social and Economic Networks*. Princeton, NJ: Princeton University Press.

Krause U (2000) A discrete nonlinear and non-autonomous model of consensus formation. In: Elaydi S, Ladas G, Popenda J, et al. (eds) *Communications in Difference Equations: Proceedings of the Fourth International Conference on Difference Equations*. London: Gordon and Breach Science Publishers.

Lehrer K and Wagner C (1981) *Rational Consensus in Science and Society: A Philosophical and Mathematical Study*. Dordrecht: Reidel.

Lorenz J (2005) A stabilization theorem for dynamics of continuous opinions. *Physica A* 355: 217–23.

Riegler A and Douven I (2009) Extending the Hegselmann-Krause model III: From single beliefs to complex belief states. *Episteme* 6(2): 145–63.

Winkler RL (1968) The consensus of subjective probability distributions. *Management Science* 15(2): B61–B75.

Zollman KJS (2007) The communication structure of epistemic communities. *Philosophy of Science* 74(5): 574–87.

Zollman KJS (2009) Optimal publishing strategies. *Episteme* 6(2): 185–99.

Zollman KJS (2010a) The epistemic benefit of transient diversity. *Erkenntnis* 72(1): 17–35.

Zollman KJS (2010b) Social structure and the effects of conformity. *Synthese* 172(3): 317–40.

About the author

Kevin JS. Zollman is Assistant Professor in the philosophy department at Carnegie Mellon University. His research focuses on social epistemology and the mathematical modeling of social behavior both in humans and animals.

Downloaded from ppe.sagepub.com at CARNEGIE MELLON UNIV LIBRARY on May 10, 2012