The Coevolution of Beliefs and Networks

Jasmina Arifovic
Simon Fraser University
arifovic@sfu.ca

B. Curtis Eaton
University of Calgary
eaton@ucalgary.ca

Graeme Walker
Simon Fraser University
gmw1@sfu.ca

First version: July 24 2014
Revised: July 20 2015

Abstract

Social psychologists have shown that people experience cognitive dissonance when two or more of their cognitions diverge, and that they actively manage the dissonance. With this in mind, we develop a model of social learning in networks to understand the coevolution of beliefs and networks. We focus on beliefs concerning an objective phenomenon. Initial beliefs are based on noisy, private and unbiased information. Because the information is noisy, initial beliefs differ, creating dissonance. In our model, behavior is motivated by a desire to minimize this dissonance. In many circumstances this behavior adversely affects the efficiency of social learning, such that in equilibrium the mean aggregate belief is biased and there is significant variation of beliefs across the population. The parameterizations of our model that result in the most inefficient learning produce a fractionalized network structure in which there are a number of distinct groups: within any group all beliefs are identical; beliefs differ from group to group, sometimes greatly; there is no intergroup interaction. Since dissonance minimizing behavior is apparently a deeply rooted feature of humans, we are led to ask: What policies could improve the situation? Our results suggest that policies that improve the availability of objective information and/or increase the size of networks enhance efficiency of social learning. On the other hand, anything that makes changing networks more attractive as a dissonance minimizing strategy has the opposite effect.

1 Introduction

Two facts motivate this paper: beliefs regarding many objective issues vary significantly over individuals in the same society; within personal networks the beliefs on such issues tend to be similar, often identical.[1] These two facts suggest that beliefs are to a significant degree socially constructed as opposed to being based solely on objective information.[2] It is also the case that the composition of personal networks is determined in part by beliefs. If the beliefs of people in one's current network are quite different from one's own beliefs, and if divergent beliefs create friction (or cognitive dissonance) in personal relationships, then one way of dealing with the friction (or decreasing the dissonance) is to develop a new network composed of people with beliefs similar to one's own. Consistent with these observations, our model is one in which beliefs and personal networks coevolve.

Kandel (1978) was the first to carefully document coevolution of this sort. Using longitudinal data on adolescent friendships from five New York State high schools during the 1971-1972 academic year, she assessed the roles of socialization and friend selection on specific behaviors and attitudes. Of roughly 1000 friendship pairs identified at the beginning of the school year, two-thirds remained at the end of the year. These friendship pairs tended to be those in which initial behaviors and attitudes were initially similar, and further they became more similar over the year. The friendship pairs that dissolved by the end of the year tended to be those in which initial behaviors and attitudes were dissimilar. As well, the initial behaviors and attitudes of people in newly formed friendship pairs tended to be similar and they became more similar over the course of the year.

It is our view that cognitive dissonance and efforts to manage it are the key to understanding the social forces that drive the coevolution of beliefs and personal networks. Carsey and Layman (2006) provide a nice illustration of how this process works. In an American context, they ask whether one's party affiliation determines one's preferences on issues, or whether one's preferences determine one's party affiliation. Using the 1992, 1994, and 1996 National Election Surveys, they construct a panel of voters and track their party affiliation, their attitudes on issues, including the relative importance of different issues to the voter, and their awareness of party differences on issues. They discover that when there is no significant difference between the individual's attitude and the position of her current party on issues of major importance to the individual, the individual tends to adopt the position of her party on most issues. However, when there are significant differences on issues of

[1]This phenomena is commonly referred to as *homophily*, which Lazarsfeld et al. (1954) define as the tendency of individuals and their associates to have similar beliefs, ethnicity, age, religion, party affiliation, education, occupation, etc. McPherson et al. (2001) provide an insightful review of this literature. For a more recent and exhaustive review on homophily, see Chapter 4 in Easley and Kleinberg (2010).

[2]The controversy over Barack Obama's place of birth illustrates both phenomena. A July 2010 CNN poll found that while forty two percent of respondents were certain that he was born in the US, more than a quarter had doubts about his birth country and more than a tenth were certain that he was born outside the US. Further, among Democrats surveyed only 15 percent thought it was possible Obama was born outside the US, compared to over 40 percent of the Republicans. Since friends of Democrats tend to be Democrats and those of Republicans to be Republicans, this data supports the second phenomena. Berelson et al. (1954) were the first to document homophily with respect to party affiiliation in the US.

major importance, the individual tends to change party affiliation by choosing to affiliate with the party whose positions on the important issues most closely reflect the individual's own attitudes.

This nicely illustrates dissonance minimizing behavior. The divergences of one's own attitudes on active issues from those of one's current party create uncomfortable cognitive dissonance. The individual can reduce that dissonance either by changing her own attitudes to conform with the positions of her current party, or by changing her party affiliation. Individuals choose the latter when there are significant differences with respect to the issues of major importance to the individual, and the former when there are no such differences.

Our objective is to develop a deeper understanding of the interplay between objective information and dissonance minimizing behaviors in the dynamic formation, or coevolution, of beliefs and personal networks.[3] We develop a model of network structure and beliefs of individuals on an objective issue. Initial beliefs are determined by noisy, private information and initial networks are randomly determined. In each period, there is a conversation concerning beliefs in which utterances are chosen to minimize dissonance. Subsequently, people adjust their beliefs and networks in light of the utterances of their associates, and again the objective is dissonance minimization. This process is repeated until beliefs and networks converge. In the resulting equilibrium, there are clusters of people with identical beliefs and the network connections of all people in any cluster are exclusively with other people in the same cluster, so in equilibrium, beliefs are perfectly homophilous. Across clusters, beliefs may vary, sometimes greatly. A large number of interesting questions can be articulated in the model. We use simulation techniques to explore these questions because the model is intractable using standard techniques.

In our model, behavior is motivated by a desire to minimize cognitive dissonance. In this, we are following the lead of Festinger (1957). Elliot and Devine (1994) introduce their experimental paper confirming this approach in the following way: "As presented in his classic monograph, Festinger (1957), cognitive dissonance theory is fundamentally motivational in nature. Festinger posited that the perception of an inconsistency among an individual's cognitions generates a negative interpersonal state (dissonance), which motivates the individual to seek and implement a strategy to alleviate this aversive state." Further, the experiments reported in Elliot and Devine (1994) provide support for the sequential structure in our model. First they create dissonance in their subjects by giving them an opportunity to advance the policy discussion concerning a proposed tuition increase. Subjects can choose to write an essay outlining the pros or the cons of the proposal, but they are informed that what is really needed is an essay in support of the proposal because the reasons for opposing it have already been thoroughly surveyed. This creates dissonance because in prior screening subjects have revealed that they strongly oppose any increase in tuition. Hence, they can advance the policy discussion only by voluntarily choosing to write an essay in support of a

[3] Akerlof and Dickens (1982) and Rabin (1994) incorporate cognitive dissonance into rational choice models and show how it can lead people to underestimate the likelihood of a bad outcome happening to them, and how it can lead people to consume too much of an immoral good, respectively. Neither of these papers investigates how cognitive dissonance influences social learning and/or networks of people, the focus of this paper.

position they strongly oppose. Subjects manage this dissonance by actually writing an essay in support of the proposal. This act, of course, creates more dissonance since they have now expressed an opinion they do not hold. Subsequently, when given a chance to revise their position on the tuition issue, they manage this dissonance by changing their position.

The foundational model for social learning in networks is a model of what has come to be known as *naïve learning* developed in Degroot (1974).[4] The naïve learning model employs an iterative process of belief updating, or social learning, wherein the beliefs of individuals are updated using a weighted average of their own current belief and the current beliefs of their network associates. Although beliefs evolve, networks are by assumption fixed. Recent applications of the model have investigated the conditions under which naïve learning is efficient in the sense that beliefs converge to the true underlying state of the world. In particular, see Golub and Jackson (2010) on the wisdom of crowds.

Our model is similar to the naïve learning model in that it involves an iterative updating process, but it differs from that model in two important ways. The first concerns the way in which people learn about the beliefs of others. The naïve learning model simply assumes that agents truthfully report their current beliefs. Given that beliefs are private information, this assumption is problematic. It would seem useful, even necessary, to ask what motivates people to truthfully reveal their private information. A possible answer is that truthful revelation is motivated by a desire to learn. While this desire clearly could lead people to make use of the truthful revelations of others, perhaps in the way posited in the naïve learning model, truthful revelation of one's own belief on some issue clearly does not in itself improve one's own learning about that issue.

For this reason we reject the truthful revelation assumption and instead begin with the question: What motivates people in the social conversation in which they make pronouncements concerning what they believe? As we thought and read about this question we were led to the work of social psychologists on what is sometimes called *audience tuning*, the tendency of communicators to "take the audience's background knowledge and opinions into account and spontaneously tune their messages to be congruent with them ..." (Higgins (1999), p.33-34)). Research on this pervasive phenomenon goes back at least as far Solomon Asch's pioneering work, Opinions and Social Pressure (Asch (1955). We model this audience tuning behavior as an optimal response to the cognitive dissonance arising in the conversation on beliefs from the divergence of one's own belief from one's own pronouncement about that belief and the divergences of one's own pronouncement from the pronouncements of others.

Having learned something about the beliefs of others via the conversation on beliefs, people still experience cognitive dissonance since not everyone professes to believe the same thing. To mange or minimize this dissonance, people revise their beliefs. We show that a person's dissonance minimizing revised belief is identical to the person's equilibrium pronouncement in the conversation on beliefs. That is, we prove a *saying is believing* result for our model. This is of some interest, given that social psychologists observe this sort of behav-

[4]See, for example, Golub and Jackson (2010), DeMarzo et al. (2003), Acemoglu, Bimpikis and Ozdaglar (2010), Acemoglu, Ozdaglar and ParandehGheibi (2010) and Bala and Goyal (1998).

ior in the laboratory. See Higgins (1999) for a discussion of audience tuning and the saying is believing effect. Although the details are quite different, as regards the results of the iterated revision of beliefs, the naïve learning model and our dissonance minimizing model deliver similar results; in particular, if networks are fixed and there is sufficient interconnectivity in them, learning is efficient.

But there is yet more to the dissonance minimizing story. If the approach is accepted, it is clearly inappropriate to assume that networks are fixed, since people can manage their dissonance by altering their networks, either eliminating or adding links with people depending on whether their beliefs are dissimilar or similar to their own. This is the second important way in which our model differs from the naïve learning model, in which networks are, by assumption, fixed. Notice that the dissonance minimizing approach leads quite naturally to a model in which beliefs and networks coevolve.

Of all the studies using a naive learning approach to social learning in networks, Golub and Jackson (2012) have motives most similar to ours. They too are interested in the phenomena of homophily and how it affects social learning. Rather than providing an explanation for what drives homophily, they introduce it exogenously by allowing nodes to be of different types and increasing the likelihood that nodes of the same type will be connected. This results in networks where nodes are clustered together by types. Unlike our model, however, the networks they analyze never have disconnected subgraphs - there are always some connections between nodes of different types, and as a result consensus is always reached. Their concern is how the degree of homophily affects the speed at which information spreads.[5]

Economists have investigated the formation and evolution of networks using two distinct models. The first, developed by Bala and Goyal (2000), considers the decisions individuals make in forming network connections through the lens of cost-benefit analysis. They model the formation of networks as a non-cooperative game and perform equilibrium analysis on the different structures of networks formed. Jackson and Watts (2002) use a similar model of endogenous network formation to investigate network properties as well as coordination in games. The focus of these studies is on how network structure affects decisions. None of these studies focus on learning.

The second type of model of network formation used by economists is one borrowed from graph theory. This model is not economic in nature in that there are no individuals making decisions, instead networks are formed probabilistically. This model has been used, successfully, to capture some of the empirical regularities of social networks (see Jackson and Rogers (2007) as a great example). This model also does not incorporate learning, as the focus of analysis is purely on network properties.

The way in which we model network formation and evolution is similar in certain respects to these models. Although we do not model the costs and benefits of making a connection

[5]Pan (2012) extends the DeGroot model by allowing the trust matrix - capturing the weights agents place on the beliefs of others - to be time-varying based on similarity of beliefs. This extension to the naïve learning model results in a consensus belief with a trust matrix that puts equal weights on all connections In other words, Pan's model converges to a particular parameterization of the DeGroot model. This parameterization coincides with our comparison with DeGroot when $beta = 1$ (i.e. equal weights placed on each associate). Importantly, like DeGroot, the network structure in Pan's model does not change.

explicitly, as do Bala and Goyal (2000), the motivation in our model for changing networks is similar in spirit. Moreover, the initial random network structure in our model is similar to Jackson and Rogers (2007).

The work on network evolution that is closest in spirit to ours comes from Schelling (1978). Without providing a formal model, Schelling (1978) discusses the evolution of social organization as being driven by homophilic preferences. As far as we are aware there is no literature in economics on the coevolution of networks and beliefs.

Holme and Newman (2006) do explore a simple simulation model of the coevolution of beliefs and networks. There are N people, each of whom holds one of G opinions, and there are M pairs of people who are linked to each other. The model is initialized by randomly assigning opinions to every person and by randomly constructing M links between pairs of people. Opinions and networks evolve through the following algorithm. In every period, a person whose opinion we will call g is randomly picked. If she is not linked to another person, nothing changes. If she is linked to one or more people, with probability Φ a randomly chosen link is dissolved and replaced with a link to a person chosen randomly from the set of people who hold opinion g, and with probability $1 - \Phi$ her opinion is changed to coincide with the opinion of one of the people with whom she is linked. Eventually, the population consists of a number of groups of people, and within each group every person holds the same opinion and no one is linked to a person outside the group. The qualitative nature of their results are exactly what one would expect. If Φ is large, changes in links are common and changes in opinion are rare so when all the dust settles, there are a larger number of groups and a large number of different opinions. On the other hand, when Φ is small, changes in links are rare and changes in opinion are common, so when all the dust settles, there are a small number of groups and a small number of different opinions. In the limit as Φ approaches 1, there is one very large group and possibly a number of very small groups. This pattern is suggestive of some of the results in the current paper.

The rest of our paper is organized as follows. In section 2, we describe our model. In section 3, we illustrate the coevolution of networks beliefs. In Section 4, we conduct comparative analyses that illuminates the way in which the key exogenous parameters affect beliefs and networks. In Section 5, we conclude.

2 The Model

We want to explore the interaction of *social processes driven by the desire to minimize dissonance* and *noisy objective private information* as they affect the personal networks and beliefs of individuals with respect to an issue where *objective truth* is a meaningful notion. This requires that we be able to identify the objective truth and to generate noisy private information. Hence, we focus on beliefs regarding stochastic phenomena. Specifically, we model beliefs regarding the mean of a beta distribution. The objective truth is just the mean of the relevant beta distribution, and noisy private information is generated by random draws from that distribution.

The beta distribution, $f(x : A, B)$, has two shape parameters, A and B, and $x \in [0, 1]$.

Depending on the values of the shape parameters, a wide array of uni and bimodal distributions are possible. The mean and variance of the beta distribution are $\frac{A}{A+B}$ and $\frac{AB}{(A+B)^2(1+A+B)}$ respectively.

2.1 Initial Conditions

There is a population of size P. In period 1, every person is given a randomly generated network of $N << P$ people, and Q private realizations of the beta distribution. P, N, and Q are positive integers.

Person i's initial network is described by a vector with P elements,

$$\mathbf{I}_i^1 = (I_{i1}^1, I_{i2}^1, ..., I_{iP}^1) \tag{1}$$

each of which is either 0 or 1. $I_{ik}^1 = 1$ if person k is a member of i's network, and $I_{ik}^1 = 0$ if she is not. $I_{ii}^1 = 0$, indicating that person i is not a member of her own network. $\mathbf{I}_i^1$ is created by randomly choosing N people from the set $\{1, 2, ..., i-1, i+1, ..., P\}$. Notice that initial network connections are not necessarily symmetric – if person n is in person m's network, person m may or may not be in person n's network.[6]

We denote the entire set of individual networks in period 1 by $\mathbf{I}^1$, where

$$\mathbf{I}^1 = (\mathbf{I}_1^1, \mathbf{I}_2^1, ..., \mathbf{I}_P^1) \tag{2}$$

Person i's private information regarding the relevant issue is $(x_{i1}, x_{i2}, ..., x_{iQ})$. Person i uses this information to form an initial belief, b_i^1, regarding the mean of $f(x; A, B)$ as follows,

$$b_i^1 = \frac{\sum_{q=1}^{Q} x_{iq}}{Q} \tag{3}$$

Her initial belief is just the arithmetic mean of her observations. The entire set of initial beliefs is

$$\mathbf{b}^1 = (b_1^1, b_2^1, ..., b_P^1) \tag{4}$$

There is another way of thinking about initial beliefs. Dissonance arises when two or more cognitions of an individual regarding the same phenomenon differ. In our model dissonance is inevitable, and we can think of the individual as choosing her initial belief to minimize the dissonance associated with the divergence of what she observes from what she believes. If we model the dissonance associated with the divergence of an observation from the belief regarding the expected value of an observation as the square of the divergence, we get the following expression for dissonance

$$\sum_{q=1}^{Q} (|b_i^1 - x_{iq}|)^2 \tag{5}$$

[6] We are using a non-cooperative approach to model coevolution. That is, people choose the things that are under their control to achieve their objectives. This requires that network connections be possibly asymmetric, for the reason that each of us can choose the people whose beliefs influence our own belief, but we cannot choose the people who will be influenced by our own belief. This approach, of course, does not preclude symmetric network connections, but it would be a mistake to require that they be symmetric.

The belief b_i^1 that minimizes dissonance is the arithmetic mean of the Q observations.

The initial state of the system is completely described by the pair $(\mathbf{I}^1, \mathbf{b}^1)$.

In any period $t > 1$, both beliefs and networks are modified through a social process that is driven by dissonance minimization. A brief sketch of this process is useful – details are provided below. Going into period t, the state of the system is $(\mathbf{I}^{t-1}, \mathbf{b}^{t-1})$. In period t, there is a *conversation* that allows people to get a fix on the current beliefs of their associates. Then, people use the information that comes out of the conversation to *revise their beliefs*. Finally, people are given the opportunity to *revise their networks* by swapping one or more people who are currently in their networks for people who are not. The utterances in these conversations and the processes by which beliefs and networks are revised are all driven by dissonance minimization. The end product is a new state of the system, $(\mathbf{I}^t, \mathbf{b}^t)$. Notice that at the end of any period, the influence of all private information and all previous social interaction is buried in the state description $(\mathbf{I}^t, \mathbf{b}^t)$.

2.2 The Conversation

Clearly, if people are to revise their beliefs in light of the beliefs of their associates, there has to be some mechanism by which they learn about the their associates' beliefs. In the literature initiated by DeGroot (1974), it is simply assumed that every person has access to the current beliefs of all associates. As discussed below, this assumption is inconsistent with the psychological literature on *message tuning*, and in fact inconsistent with the usual view in economics that actions are chosen to maximize the utility of the actor.[7] In our model of the conversation, people choose their message to minimize dissonance.

People enter the conversation with a current belief. In period $t \geq 2$ person i's current belief is b_i^{t-1}. In the conversation, people make utterances or announcements regarding their own beliefs. In period t, person i's utterance is a_i^t. People pay attention only to the utterances of people in their network.

In this conversation, there are two sources of dissonance. Dissonance arises when a person's own utterance diverges from her own current belief, and the magnitude of the dissonance on this account is $\alpha(|a_i^t - b_i^{t-1}|)^2$, where $\alpha > 0$ is the *private* dissonance parameter.[8] We use the adjective private because the dissonance that arises when a person says one thing while actually believing another is internal or private. Dissonance also arises when a person's utterance differs from that of one of her associates, person k for concreteness, and the magnitude of the dissonance is $\beta(|a_i^t - a_k^t|)^2$, where $\beta > 0$ is the *social* dissonance parameter. We assume that $\alpha \geq \beta$.

[7]Kuran (1997) is a noteworthy exception. The distinguishing feature of this important book is the hypothesis that people choose what they say to achieve the consequences and impressions that serve them best.

[8]Notice that having dissonance proportional to the *squared* difference is merely an assumption we make. In Chapter 5 of his dissertation, Walker (2015) relaxes this assumption by replacing the exponent 2 with an exponent γ that is greater than one. He explores various values of γ, both larger and smaller than 2. As this parameter nears one, the dissonance function is approximately linear. Such dissonance functions result in beliefs converging to the median utterance in one's network where equilibrium beliefs can be completely polarized and networks are split in two along these polarized beliefs.

The aggregate dissonance of person i in the period t conversation is then

$$D_i^t(a_i^t) = \alpha(|a_i^t - b_i^{t-1}|)^2 + \beta \sum_{k=1}^{P} I_{ik}^{t-1}(|a_i^t - a_k^t|)^2 \tag{6}$$

Person i's objective is to minimize her dissonance, and her instrument in the conversation is her own utterance. Accordingly, person i chooses a_i^t to minimize $D_i^t(a_i^t)$. $D_i^t(a_i^t)$ is strictly convex in a_i^t, so person i's best response is implicitly defined by $\partial D_i^t(a_i^t)/\partial a_i^t = 0$. This condition can be written as

$$\alpha(a_i^t - b_i^{t-1}) + \beta \sum_{k=1}^{P} I_{ik}^{t-1}(a_i^t - a_k^t) = 0 \tag{7}$$

Although we used two dissonance parameters, α and β, one of them is redundant, since in the dissonance minimizing game, the relative magnitude of the two parameters is all that matters. If $\widehat{a}_i^t$ satisfies condition (7), given $\alpha = \widehat{\alpha}$ and $\beta = \widehat{\beta}$, it is the case that $\widehat{a}_i^t$ also satisfies condition (7), given $\alpha = 1$ and $\beta = \widehat{\beta}/\widehat{\alpha}$. Accordingly, we fix $\alpha = 1$, and we are left with just one dissonance parameter, $\beta \leq 1$.

The Nash equilibrium of the dissonance minimization game, $\mathbf{a}^{t^*} = (a_1^{t^*}, a_2^{t^*}, ..., a_P^{t^*})$, is then the solution of the following system of first order conditions:

$$(a_i^{t^*} - b_i^{t-1}) + \beta \sum_{k=1}^{P} I_{ik}^{t-1}(a_i^{t^*} - a_k^{t^*}) = 0, i = 1, P \tag{8}$$

Given that all beliefs are in interval $[0, 1]$, there is a unique solution in which all equilibrium utterances are also in interval $[0, 1]$.

Obviously, $\mathbf{a}^{t^*}$ is the equilibrium of the game in which utterances are chosen simultaneously to minimize dissonance. Clearly we cannot usefully regard the real conversation that we are attempting to model as a simultaneous move game.[9] If people were forced to choose their utterances simultaneously, there is virtually no chance they would choose utterances $\mathbf{a}^{t^*}$ – with no inkling of what others are going to say and no good way of predicting what they will say, people have no way of implementing their desire to minimize dissonance. Instead, we conceive of the real conversation within the period as a dynamic game in which people initially offer tentative utterances about their own beliefs, and then repeatedly revise their tentative utterances in light of the tentative utterances of their associates until tentative utterances converge. This raises an obvious question. Is there any reason to think that such a process would converge to $\mathbf{a}^{t^*}$? We have conducted a number of simulations of the conversation using the following algorithm: initial tentative utterances are identical to current beliefs; in a predetermined sequence, players repeatedly use their best response functions to alter their tentative utterances. In all of the simulations that we have examined, the process

[9] In Chapter 3 of Walker (2015), Walker explores an alternative model of the conversation. In the alternative model, utterances are chosen conditional on one's current beliefs and the previous utterances of one's associates. The tenor of results, but not the details, is the same as in this paper.

does converge to equilibrium utterances. We have not, however, been able to prove that this is always the case.

Our model of the conversation captures what social psychologists call *audience tuning*, crafting one's communications so that they are congruent with the opinions held by one's audience, as a dissonance minimizing behavior. Higgins (1999, p. 33-34) puts it this way:

> Your boss asks your opinion on some issue. You already have some idea of your boss's own opinion on the issue. What do you do? If you are like most people you take your boss's opinion into account when communicating your own opinion. Such tailoring and tuning of messages to suit one's audience is common in interpersonal communication. ... Indeed, this is true even at the basic level of assuring the communicator and the audience have the same reference or topic in mind for the communication. ... But beyond common reference for the communication, communicators also take the audience's background knowledge and opinions into account and spontaneously tune their messages to be congruent with them ... Moreover, such spontaneous tuning is not limited to audiences who are high in power or status ... . People tune to all sorts of audiences.

2.3 Revising Beliefs

In every period, every person uses what they learn in conversation about the beliefs of their associates to revise their own belief. We assume that people take the equilibrium utterances of their associates to be the current beliefs of their associates, which allows us to formulate a choice problem in which people choose a revised belief to minimize the dissonance associated with the divergences of their revised belief from the current beliefs of their associates and their own current belief. Importantly, the assumption that the current beliefs of associates are identical to their equilibrium utterances turns out to be correct.

Again there are two sources of dissonance. One source is associated with the difference between the person's revised belief, b_i^t, and her current belief, b_i^{t-1}, and the dissonance on this account is $(|b_i^t - b_i^{t-1}|)^2$. This is a private matter, so we use the private dissonance parameter, $\alpha = 1$. The other source is associated with the difference between the person's revised belief, b_i^t, and her perception of the current belief of each of her associates – for associate k, the dissonance on this account is $\beta(|b_i^t - a_k^{t^*}|)^2$. Since this is a social matter, we use the use the social dissonance parameter, β.

Person i's aggregate dissonance is then

$$\widetilde{D}_i^t(b_i^t) = (|b_i^t - b_i^{t-1}|)^2 + \beta \sum_{k=1}^{P} I_{ik}^{t-1}(|b_i^t - a_k^{t^*}|)^2$$

The variables b_i^{t-1} and $a_k^{t^*}$ are predetermined at the point in time when person i is revising her belief, so belief revision is a decision problem, not a game. Since $\widetilde{D}_i^t(b_i^t)$ is strictly convex in b_i^t, person i's dissonance minimizing revised belief is implicitly defined by $\partial \widetilde{D}_i^t(b_i^t)/\partial b_i^t = 0$, which can be written as

$$(b_i^t - b_i^{t-1}) + \beta \sum_{k=1}^{P} I_{ik}^t (b_i^t - a_k^{t^*}) = 0 \tag{9}$$

When we compare this condition with condition (8), it is clear that the dissonance minimizing belief of person i is identical to her equilibrium utterance:

$$b_i^t = a_i^{t^*}$$

Further, her assumption that the equilibrium utterances of her associates represent their current beliefs turns out to be accurate. In short, this formulation of the belief revision problem yields a *saying is believing* result.

There is a large literature in social psychology on what is often called the *saying is believing effect* that is consistent with our belief revision result.[10] It is worth quoting Higgins (1999, p. 42) again:

> There are various kinds of social influence that will induce people to express ideas that are not their own. Across a variety of literature, including role-playing, cognitive dissonance, and social cognition, the striking finding is that communicators often end up believing what they say.

Before we turn to an examination of the way in which people revise their networks in our model, it is useful to say something about the relationship between our cognitive dissonance minimizing approach to the revision of beliefs and the naïve learning approach. The first order conditions for equilibrium utterances defined in equations (8) can be written as $\boldsymbol{a}^t \boldsymbol{A} = \boldsymbol{b}^{t-1}$, where $\boldsymbol{a}^t$ is a 1 by P vector of utterances in period t, $\boldsymbol{b}^{t-1}$ is a 1 by P vector of predetermined beliefs, and $\boldsymbol{A}$ is a coefficient matrix. Element A_{ij} of $\boldsymbol{A}$ is $1 + \beta N$ if $j = i$, 0 if person i is not in j's network, and β if person i is in i's network. Period t equilibrium utterances are given by $\boldsymbol{a}^t = \boldsymbol{b}^{t-1}\mathbf{A}^{-1}$ and period t revised beliefs by the same expresion, $\boldsymbol{b}^t = \boldsymbol{b}^{t-1}\mathbf{A}^{-1}$. In the naïve learning model period t beliefs are produced from period $t-1$ beliefs in exactly the same way if the trust matrix is $\boldsymbol{T} = \mathbf{A}^{-1}$. In this sense, there is an equivalence between the dissonance minimizing approach and the naïve learning approach. That is, if we confine out attention to the way in which period $t-1$ beliefs are transformed into period t beliefs, given any permissible $\boldsymbol{A}$ matrix there is trust matrix $\boldsymbol{T} = \mathbf{A}^{-1}$ that transforms beliefs in exactly the same way.

However, if the dissonance minimizing approach is accepted, this equivalence is clearly not the end of the story because people can manage their dissonance not only by changing their beliefs, but also by altering their networks. In other words, if the dissonance minimizing approach is accepted, the fixed network assumption of the naïve learning model must be rejected.

[10] The *saying is believing* effect can be traced back to Festinger and Carlsmith (1959), which offers "a test of some hypotheses generated by Festinger's theory of cognitive dissonance, viz., that *if a person is induced to do or say something which is contrary to his private opinion, there will be a tendency for him to change his opinion to bring it into correspondence what he has done or said.*"

2.4 Revising Networks

In every period, after beliefs have been revised, every person i is given the opportunity to reduce her dissonance by revising her network. Person i serially meets W potential new associates, randomly chosen from the set of people not currently in her network, and in each case is given access to their current belief.[11] W is a non-negative integer. For each potential new associate, person i considers replacing the current associate whose belief is furthermost from her own with the potential new associate, and does so if and only if the swap reduces her dissonance by at least s, where s is a small real number that can be thought of as a *switching cost*. In the simulations reported here we used $s = 0.000001$.[12]

Our dissonance minimizing approach to network revision suggests a perhaps novel interpretation of the large body of literature in social psychology on confirmation bias. Nickerson (1998, p. 175) provides a survey of this literature, and he starts with the following definition:

> Confirmation bias, as the term is typically used in the psychological literature, connotes the seeking or interpreting of evidence in ways that are partial to existing beliefs, expectations, or a hypothesis in hand.

Clearly, network revision as we model it produces a confirmation bias, since over time people quit listening to others whose beliefs differ significantly from their own. The novel interpretation is this: we would suggest that the confirmation bias phenomenon is driven, at least in part, by a desire to minimize dissonance. Further, if this suggestion is accepted, it is apparent that the dissonance minimizing approach to coevolution provides an analytical synthesis of the literatures in social psychology on audience tuning, saying is believing, and confirmation bias.

2.5 Convergence

If $t \geq 20$, we check for convergence to an equilibrium. We use two criteria. Convergence requires (i) that in periods $t-20$ through t there be no change in any person's network, and (ii) that the sum of $|b_i^t - b_i^{t-1}|$ over periods $t-20$ through t and all P individuals be less than 0.001. If both criteria are satisfied, we deem the simulation to have converged and we report the pair $(\mathbf{I}^{t^*}, \mathbf{b}^{t^*})$ as the equilibrium of the simulation.

2.6 Outcome Indices

It is simply not feasible to report all of the data that our simulations generate. We have chosen to report five indices that we think capture most of the interesting results.

[11] The obvious alternative to the completely random process for choosing potential new associates would be to have the condition that potential new associates be drawn from the set of people who are connected to one's own current associates. Chapter 4 in Walker (2015) extends our model in this direction and shows that having a more directed network revision process leads to equilibria with more polarized beliefs and more fragmented networks.

[12] Huston and Levinger (1978) provide experimental evidence that is consistent with our network revision algorithm.

Like DeGroot's, ours is a model of the social aggregation of private information, and we report three indices, $ENLITE$, $STDEV$, and $BIAS$, that give some sense of how efficiently information is aggregated. Given the private information that individuals have in period 1, the best estimate of the mean of the beta distribution is

$$\bar{b}^1 = \frac{\sum_{i=1}^{P} b_i^1}{P}$$

Initial ignorance, IG^1, is then $\sum_{i=1}^{P} |b_i^1 - \bar{b}^1|$. Similarly, ignorance in equilibrium, IG^E, is $\sum_{i=1}^{P} |b_i^E - \bar{b}^1|$, where b_i^E is person i's equilibrium belief. $ENLITE$, short for enlightenment, is defined as follows:

$$ENLITE = \frac{IG^1 - IG^E}{IG^1}$$

The maximum value of $ENLITE$ is 1 which results when each person has learned the correct mean of the beta distribution. We also report the standard deviation of equilibrium beliefs, $STDEV$, and an index of the bias in equilibrium beliefs:

$$BIAS = \bar{b}^E - \bar{b}^1$$

where $\bar{b}^E$ is the mean equilibrium belief.

We report two indices that summarize key features of the structure of equilibrium networks, $SECTS$, and CON. $SECTS$ is the number of disconnected subgraphs in the equilibrium network structure. The group of individuals who form a particular disconnected subgraph, or sect, have network connections (direct or indirect) only with other members of that sect. Obviously, the value of $SECTS$ is at least 1, and could be as large as P/N.

Consider any network structure with n sects, and let s_i denote the number of people in sect i, where $i \in \{1, ..., I\}$. A representative person in sect i is connected, directly or indirectly, to the other $s_i - 1$ people in sect i, and she has no connection with the other $P - s_i$ people in the population. The fraction $\frac{s_i - 1}{P}$ is an index of the connectivity of any person in sect i with the entire population of P people. CON is the mean connectivity of the entire population:

$$CON = \sum_{i=1}^{n} (\frac{s_i}{P})(\frac{s_i - 1}{P - 1}) \tag{10}$$

CON is bounded below by $\frac{N-1}{P}$ and above by 1. Notice that CON is the answer to the following question: If, in equilibrium, the belief of a randomly selected person is exogenously altered, what is the expected proportion of the population that will desire to alter their own belief?

2.7 Parameters

Our model involves a total of seven parameters: the shape parameters of the beta distribution (A and B); the size of the population (P); the *quantity of objective information* available to individuals (Q); the *sensitivity to cognitive dissonance* (β); the *size of personal networks* (N); the *frequency of opportunities to meet new potential associates* (W).

As it turns out, if P is moderately large relative to N, variations in P have no interesting effects on equilibrium beliefs and networks, and in our results section P is constant.

Four of the model parameters, Q, N, β and W either capture directly or act as proxies for important social variables and processes, and in the next section we pay a lot of attention to the implications of variations in these parameters. The importance of objective information is obvious, so understanding the ways in which the equilibrium responds to variation in Q is clearly of interest. The implications of variations in sensitivity to dissonance, parameter β, are also of great interest to us, since the desire to minimize dissonance drives the choices of individuals make in out model.

Network size, N, directly captures the extent of social interaction in the model. In the background, however, are the costs and benefits of social interaction that are unrelated to the dissonance considerations that are central to our analysis. Since in equilibrium dissonance vanishes, network size is ultimately driven by these costs and benefits. Indeed, in earlier work on this model, N was endogenous, and driven by precisely these benefits and costs. We found, however, that in the interests of clarity it was best to make N exogenous, pushing the utility calculus that drives network size into the background. Recent developments in communication technology that are familiar to all have significantly reduced the costs of network interaction, so comparative dynamic exercises in which N increases are of some interest.

W can be thought of as a proxy for the frequency with which people have the opportunity to amend their personal networks. The frequency will be high in situations where people have many opportunities to interact with others who are not currently close to them. Thus W can be seen as a proxy for the openness of society. The more open is a society, the easier it is to effect dissonance reduction by altering one's personal network, so the extent to which dissonance minimization is achieved by changes in network composition as opposed to changes in beliefs is directly related to the magnitude of W.

3 Results for a Bimodal Distribution

Individuals form beliefs regarding the mean of a beta distribution. Initial beliefs are based on random, independent drawings from the distribution. Then, driven by dissonance minimizing behaviors, beliefs and networks coevolve. The efficiency of the social aggregation of beliefs that occurs in our model is dependent on the properties of the underlying beta distribution, and in particular on whether the distribution is bimodal or unimodal. In this section we report results for a radically bimodal beta distribution. The mean, or objective truth, is .6. The distribution closely approximates a Bernoulli distribution in which the probability of a 1 is .6 and the probability of a 0 is .4, and in this sense this distribution is radically bimodal.[13]

[13] The shape parameters are A=0.003 and B=0.002. The mean and standard deviation of the distribton are .6 and 0.4886, respectively. The probability that a random draw is less than .02 is approximately 0.4,

3.1 Fixed Networks ($W = 0$)

To establish a point of reference, we start with results for the case in which beliefs evolve, but networks do not (i.e. $W = 0$). Keeping the network structure fixed allows us to compare the results of our model to DeGroot's naive learning model in which the network structure is also fixed. In Figures 1, 2 and 3 only, we present results for both our model and a DeGroot model in which the weight T_{ij} that person i puts on person j's belief in the updating process is 0 if person $j \neq i$ is not in i's network, is $\frac{\beta}{1+N\beta}$ if person $j \neq i$ is in i's network, and is $\frac{1}{1+N\beta}$ if $j = i$.

The evolution of the beliefs for a typical simulation is illustrated in Figure 1. In the simulation, $W = 0$, $Q = 1$, $N = 4$, $\beta = .2$ and $P = 200$. Results for our model are presented in the fourels on the left side of the figure and those for the DeGroot model in the four panels on the right side of the figure. The same seed is used in both simulations, and hence initial conditions are identical. Equilibrium beliefs are pictured at the bottom of the figure ($t = converged$), initial beliefs at the top ($t = 1$), and beliefs for two intermediate points in the simulation ($t = 5$ and $t = 10$) in the second and third rows. Along the left edge of each panel, people are ordered by their equilibrium beliefs, from highest to lowest moving upward from the origin. Since there are 200 people in the simulation, in each panel there are 200 rows. Row 1 (the lowest row in the figure) is for the person with the highest equilibrium belief, and row 200 (the highest row) is for the person with the lowest equilibrium belief. A person's belief is given by the length of the bar for her row. Initial beliefs are identical in the two models – approximately 120 people have an initial belief very close to 1, and approximately 80 have an initial belief very close to 0. For both models, over the course of the simulation beliefs of all people converge to the same value, approximately .6. The only noticeable difference between the models is that beliefs converge more rapidly in our model than in the DeGroot model. [14]

In Figure 2 we look at the same simulations in a different way. We first assign people to groups, a high belief or HB group if their initial belief is close to 1 and a low belief or LB group if their initial belief is close to 0. In the figure we report for each period the mean belief for the HB group, the LB group, and the entire group. Throughout the simulation, the mean belief for the entire group is approximately .6. The story is quite different for the HB and LB groups; initial beliefs in the two groups are very close to 1 and 0, respectively, and as the simulation proceeds the mean beliefs of the two groups converge to approximately 0.6. Again the only noticeable difference between our model and DeGroot's is the speed of convergence. With fixed networks, the two models deliver the same equilibrium results.[15]

The results reported in Figure 3 refer to our model's belief-based indices. The figure

and the probability that it is greater than .98 is approximately 0.6.

[14] In any period, utterances in the conversation that take place in the AEW model are the equilibrium to the game in which all people choose an utterance to minimize dissonance. The alternative would seem to be to choose utterances as best responses to current beliefs. The alternative would clearly slow the convergence process and it is closer to the iterative process in the naive learning model. Hence, the difference in convergence speed between the AEW and naive learning models seems to be attributable to the fact that AEW uses an equilibrium model of the conversation.

[15] Walker (2014) presents results for a systematic comparison of the two models that confirm this result.

reports averages of $ENLITE$, $STDEV$, and $BIAS$ over 1000 simulations, each with different randomly chosen initial conditions. In both sets of simulations, $N = 4$, $Q = 1$, and $W = 0$. Results for a number of values of β in interval [0,1] are reported.[16] The story told in Figure 3 regarding beliefs is one of efficient social aggregation for both our model and the naive learning model. The ignorance that characterizes initial beliefs is almost totally eliminated (the average values for $ENLITE$ are close to 0.97), beliefs converge ($STDEV$ is very close to 0), and they converge to the objective truth ($BIAS$ is very close to 0). The take-away message is clear: *when networks are chosen randomly and private information is generated by random independent draws from the beta distribution, so long as* $W = 0$ *our model aggregates private information efficiently.*[17]

3.2 Coevolution and Bias Illustrated ($W > 0$)

Once networks as well as beliefs are allowed to evolve, the sharp picture of efficient social aggregation of information we painted in the preceding subsection starts to blur. Exactly how blurry depends on a number of factors, as detailed in the next subsection where we present systematic results. In this subsection, our first objective is to illustrate the process of coevolution and the way in which it blurs the efficient aggregation picture, and our second is to develop an intuitive understanding of the forces that distort the efficient aggregation picture. Here, as in everything that follows, we focus exclusively on our own model.

Coevolution is illustrated in Figure 4, where we present results for a modified version of the simulation we considered in Figure 1. The one modification is that W is 1 as opposed to 0. For purposes of comparison, in the last column of Figure 4 we show the evolution of beliefs for the $W = 0$ case (that we first saw in Figure 1) and in the middle column the evolution of beliefs for the $W = 1$ case. With fixed networks ($W = 0$), beliefs of all people converge to a common value, whereas when networks and beliefs coevolve ($W = 1$) they do not converge – in fact, in equilibrium there are two distinct beliefs. Although it is not immediately clear from Figure 4, the mean belief in the $W = 1$ case is substantially higher than the mean belief in the $W = 0$ case.

In the first column of Figure 4, the evolution of networks in the $W = 1$ simulation is illustrated. Along the left edge of each panel in the first column, people are ordered by their equilibrium beliefs, from highest to lowest moving upward from the origin, and along the bottom edge of each panel, people are also ordered by their equilibrium beliefs, from lowest to highest moving rightward from the origin. There are 200 rows in each panel, and in every row there are four dots. In row j, the four dots represent the four people who are currently in person j's network. In the first period of the simulation ($t = 1$), the 800 dots are randomly scattered about the panel, because initial networks are random. As the simulation proceeds, we see the emergence of two sects, and from the second column it is apparent that the emergence of the two sects is accompanied by the convergence of beliefs

[16]Specifically, β took on a value from the set {.01, .05, .1, .2, .3, .4, .5, .6, .8, 1 }.

[17]Based on results from a number of simulations not reported here, the efficient aggregation story told in Figure 3 is not dependent on the values of Q and N.

of people in each sect to a common value. The network revision decisions of people drive the emergence of the two groups, because a person revises her network only when she has the opportunity to substitute a new associate whose belief is closer to her own than is the belief of the person in her current network who creates the most dissonance.

In Figure 5 we illustrate the prominent bias in equilibrium beliefs that arises in this particular simulation. For purposes of comparison we report results for both the $W = 0$ case, in the left panel of the figure, and the $W = 1$ case, in the right panel. Similar to Figure 2, this figure sorts people into a low belief (LB) group and a high belief (HB) group. In the figure we report three things for the first 90 periods of the simulation: the mean beliefs in the HB and LB groups and the mean belief for the entire population. In both panels of the figure, the mean belief for the HB group starts out very close to 1 and that for the LB group very close to 0. Not surprisingly, throughout the course of both simulations the mean HB belief falls and the mean LB belief rises. Notice however, that in the $W = 1$ case, the mean beliefs of the two groups fail to converge to a common values, as they do in $W = 0$ case. Notice also that a definite bias is seen in the evolution of the mean belief for the entire group in the $W = 1$ case. In period 1, the mean belief of the entire group is approximately .6. This initial value is, of course, the best estimate of the objective truth for this simulation, given the private information available to people – recall that the mean value of initial beliefs is identical to the mean value of all the pieces of private information. As the simulation proceeds, the mean value of beliefs for the entire group rises, eventually approaching a value that is approximately .7. Thus, in the $W = 1$ simulation, the social aggregation of information has a marked upward bias of approximately .1.

To understand the source of this upward bias, first recall that the beta distribution that is used in these simulations is radically bimodal: 60% of the mass is very close to 1, and 40% is very close to 0. In consequence, the distribution of beliefs in the first few periods of a simulation will have many more people at the top end of the distribution than at the bottom end, and further the probability that potential new associates encountered in the network revision scenario will be HB will be significantly larger than the probability that they will be LB. Thus, the probability that HB people will replace current LB associates with new HB associates is significantly larger than is the probability that LB people will replace current HB associates with new LB associates. Consequently, HB people are likely to resolve the personal dissonance created by LB associates by replacing them with HB associates, whereas LB people are likely to resolve the personal dissonance created by HB associates by revising their beliefs upward. The net effect is an upward bias in the mean final beliefs.

3.3 Systematic Results for Variations in N, Q, W and β

Here we systematically explore the way in way in which N (network size), Q (amount of private information), W (number of potential new associates), and β (sensitivity to cognitive dissonance) affect aspects of equilibrium beliefs and networks. In this analysis, we use the extreme bimodal beta distribution introduced above and a population size of 200 ($P = 200$). Results reported in Figures 6 through 10 pertain to simulations for 180 sets of parameter values.[18] For each of the 180 parameter combinations, we ran 1000 simulations, each with a different seed. Results reported in the figures are mean values of the relevant indices over the 1000 simulations. The dashed lines in the figures pertain to results for simulations in which $N = 4$, and the solid lines to simulations in which $N = 8$. Additionally, in tables in the appendix we report the mean values of the indices and the 5th and 95th percentiles of the 1000 simulations for 54 of the 180 parameter combinations.[19]

Let us first look at results for STDEV (the standard deviation of equilibrium beliefs). Results presented in Figure 6 suggest 4 intuitive results. A comparison of the dashed and solid lines in each panel reveals that N and STDEV are inversely related – the larger are networks, the smaller is the variation in equilibrium beliefs. A comparison of the slopes of the dashed and solid lines in each panel reveals that β and STDEV are inversely related – the more sensitive people are to cognitive dissonance, the smaller is the variation in equilibrium beliefs. Comparisons of the 3 panels in each row of the figure reveal that W and STDEV are directly related – the more opportunities people have to reduce dissonance by picking their associates, the larger is the variation in equilibrium beliefs. Comparisons of the 3 panels in each column of the figure reveal that Q and STDEV are inversely related – the more private information people have, the smaller is the variation in equilibrium beliefs. From the table in the appendix for STDEV we get some idea of the circumstances in which STDEV is and is not significantly different from 0, or of the circumstances in which the social aggregation of information does and does not produce consensus. For example, when $Q = 1, \beta = 0.6, N = 4$ and $W = 4$, we see that the mean value of STDEV in the 1000 simulations we conducted is 0.0882, the 5th percentile is 0.0596 and the 95th is 0.1124. In this case, the expected value of STDEV in any one simulation is significantly larger than 0. We get quite different results when $Q = 2, \beta = 0.6, N = 8$ and $W = 2$. For this parameter combination, the mean value of STDEV, the 5th percentile, and 95th percentile of the distribution of mean values are all 0 (up to 3 decimal places) – beliefs of individuals converge to a common value.

Now let us look at results for ENLITE, presented in Figure 7. ENLITE is an index of the extent to which the aggregate ignorance that is present in the initial beliefs of individuals is eliminated in the course of the coevolution that occurs in our simulations A negative value of ENLITE would indicate an increase in aggregate ignorance, a zero value no change in ignorance, and a value of 1 the complete elimination of ignorance. The first thing to note about Figure 7 is that ENLITE is always positive. This is perhaps best seen by looking at the table in the appendix for ENLITE. In the table β values are in the set $\{0.2, 0.4, 0.6\}$. The smallest mean value of ENLITE reported in the table is .5145 (associated with $Q = 1$, $N = 4$, $W = 4$, and $\beta = .2$), all other mean values are greater than .6, and most are greater than .8. The width of the (.05,.95) confidence intervals range in size from a low of about .05 to a high of about .12, so in most of the individual simulations ENLITE is well above .5. The maximum value of ENLITE is roughly .975, and this occurs for a number of different

[18] The parameter values for the 180 simulations are all possible combinations of N, Q, W and β in which $N \in \{4, 8\}$, $Q \in \{1, 2, 4\}$, $W \in \{1, 2, 4\}$, and $\beta \in \{0.01, 0.05, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.8, 1\}$.

[19] The parameter values for the 54 simulations are all possible combinations of N, Q, W and β in which $N \in \{4, 8\}$, $Q \in \{1, 2, 4\}$, $W \in \{1, 2, 4\}$, and $\beta \in \{0.2, 0.4, 0.6\}$. In each table, for each value of β there are three rows of values reported representing the 5th percentile, the mean, and the 95th percentile respectively.

parameter combinations – for all them the 5th and 95th percentiles of the distribution of ENLITE over the 1000 simulations are approximately .935 and .998.

Figure 7 suggests three intuitive comparative results. N and ENLITE are directly related – the larger are networks, the more effectively does social aggregation eliminate ignorance. W and ENLITE are inversely related – the more opportunities people have to reduce dissonance by picking their associates, the less effectively does social aggregation eliminate ignorance. β and ENLITE are directly related – the more sensitive people are to cognitive dissonance, the more effectively does social aggregation eliminate ignorance. There is no strong, systematic relationship between Q and ENLITE, nor should we expect to see one. Q has a significant, inverse effect on the absolute amount of initial ignorance, but ENLITE is an index of the extent to which initial ignorance, whatever its magnitude, is eliminated in the coevolution of beliefs and networks.

Results for CON, our index of connectivity, are presented in Figure 8. Every person is connected, either directly or indirectly, to everyone in the sect of which she is a member. For any person, connectivity is in the interval $[\frac{N-1}{P-1}, 1]$ since, at one extreme, the person could be connected only to the members of her own network, and at the other extreme to everyone else in the entire population. If there is just 1 sect in a simulation, for that simulation CON is 1, but logically there could be as many as $\frac{P}{N}$ sects, in which case CON would be $\frac{N-1}{P-1}$. The values for CON in Figure 8 and in the tables in the appendix, are of course, averages over 1000 simulations for each combination of parameters. Figure 8 suggests three intuitive comparative results. N and CON are directly related – the larger are networks, the more connected is the population in equilibrium. W and CON are inversely related – the more opportunities people have to reduce dissonance by picking their associates, the less connected is the population in equilibrium. β and CON are directly related – the more sensitive people are to cognitive dissonance, the more connected is the population in equilibrium.[20] There is no strong, systematic relationship between Q and CON. These results mirror results concerning ENLITE in the previous paragraph. Although ENLITE and CON are not perfectly correlated, the correlation is very high – the forces that eliminate ignorance through social interaction also tend to promote connectivity.

As a comparison of Figures 8 and 9 reveals, CON and SECTS are highly, negatively correlated indices. This makes good sense – if, in every simulation, the sizes of all sects were identical CON and SECTS would perfectly, negatively correlated. In fact, the sizes of sects within any simulation are not the same and sizes and number of sects changes from one simulation to another, so CON and SECTS are not perfectly correlated.

Results for BIAS are presented in Figure 10. The first thing to note is that there is a systematic upward bias in beliefs. The figure reports, for each of 180 parameter combinations, the mean bias in 1000 simulations, and this statistic is positive for every one of the 180 parameter combinations.[21] Further, as seen in the appendix table for BIAS, in the

[20] All three of these relationships seem to make intuitive sense. However, for very small values of β, the relationships do not hold. We have no good explanation of this.

[21] As explained in the previous section, the positive bias arises because the mean of the underlying beta distribution in these sumulations is .6, which exceeds .5. A negative bias woulod arise if the mean of the underlying beta distribution was less than .5.

first row of the figure, where $Q = 1$, the 5th percentile of the distribution of the statistic is also positive. Just one unambiguous comparative result is suggested by Figure 10. Q and BIAS are negatively related – the more private information people have, the smaller is the expected bias in equilibrium beliefs. A non-linear relationship between BIAS and β is evident in Figure 10 – when β is small, BIAS increases with β, and when β is large, BIAS decreases with β. A couple of observations explain this curious result: when β is very small, beliefs converge very slowly, and consequently in the coevolution, networks change radically and beliefs very little, and since the mean initial belief is unbiased, $BIAS$ is very small. At the other extreme, when β is very large, beliefs converge so rapidly that opportunities to reduce dissonance by changing networks quickly disappear, and again $BIAS$ is small. For intermediate values of β, the scenario we saw in Figures 4 and 5 occurs, and significant positive $BIAS$ results.[22]

4 Results for a Unimodal Distribution

In this section we report some results for a unimodal beta distribution. The values of the shape parameters are $A = 3$ and $B = 2$. The mean of the unimodal beta distribution is .6, as was the mean of the bimodal distribution used in the previous section. The standard deviation is, of course much smaller, 0.2 versus 0.4886. Our purpose in this section is to show that many of the qualitative features of the results reported for the bimodal distribution hold for a unimodal distribution.

In Figure 11 we report comparative results for twenty sets of parameter values, ten for each parameterization of the beta distribution. In each set for each of the beta distribution parameterizations, $N = 4, Q = 1, W = 4, P = 200$ and β is an element of the set $\{.01, .05, .1, .2, .3, .4, .5, .6, .8, 1\}$. The dashed line in each panel pertains to the radically bimodal distribution and the solid line to the unimodal distribution. The qualitative patterns are the same for the two distributions, but the social aggregation of private information is significantly more efficient for the unimodal distribution: values for $ENLITE$ are larger, and values for $STDEV$ and $BIAS$ are smaller. Except for very small values of β, results for $SECTS$ and CON are remarkably similar.

5 Conclusion

Economists typically model the formation of beliefs and networks independently. Social psychologists, however, tell us that our beliefs are dependent on the beliefs of the people with whom we associate, and at the same time, that the composition of our networks is dependent on the beliefs of others. It is in this spirit that we develop a model of the coevolution of beliefs and networks.

[22] As an exercise, we investigate the stability of our model by perturbing either the beliefs or networks for 1%, 3% or 5% of the agents in a simulation. Although these perturbations resulted in anywhere from 4% to 33% of agents updating their networks, changes to all our reported outcomes - namely STDEV, BIAS, ENLITE, SECTS and CON - are indistinguishable from zero.

Standard economic models of the evolution of beliefs in networks make two critical assumptions. People truthfully communicate their beliefs to others and communication networks do not change. In an attempt to build a model on firmer behavioral foundations, we use neither assumption.

First, we relax the assumption of truthful communication. Clearly, in order for beliefs to be communicated, there needs to be some sort of conversation. In our model of this conversation, consistent with the literature on audience tuning, people balance the dissonance that is associated with saying something that is different from their true belief, on the one hand, and saying something that is different from what their associates are saying, on the other hand. Subsequently, people revise their beliefs, in light of what they have learned in this conversation about the beliefs of their associates, and again the objective is to minimize dissonance. Revised beliefs turn out to be identical to those that were expressed in conversation. Thus our model generates a *saying is believing* result that is entirely consistent with the saying is believing phenomenon. Second, we allow networks to evolve by periodically giving people the opportunity to replace a current associate with a potential new associate. The replacement decision is also motivated by the desire to minimize dissonance, and is clearly consistent with the *confirmation bias* phenomenon. Notice that our dissonance minimizing model provides an analytical synthesis of the literatures in social psychology on audience tuning, saying is believing, and confirmation bias.

In our model, beliefs and networks evolve recursively. In the first period, people are randomly assigned a network of associates and are given a quantity of noisy objective information that they use to form an initial belief. Subsequently, there is a sequence of conversations and opportunities to alter networks that ends only when people cease to change their beliefs and networks. In this equilibrium there are groups of people, we call them sects, with identical beliefs, and the networks of everyone in a particular sect are composed of others from that sect.

From one perspective, ours is a model of the social aggregation of private information, which raises the following questions. Is this way of aggregating information efficient? In particular, how effectively does the process create informed citizens? How effectively does it reduce the variation in beliefs across the population? Does the process produce biased beliefs?

In addressing these questions, we focus on the impact of four exogenous variables in the model: the amount of objective, private information available to individuals (Q); network size (N); the sensitivity to cognitive dissonance (β); the number of opportunities individuals have in every period to replace a current associate with a new one (W).

We find a number of unambiguous results. The most striking result concerns bias in beliefs – so long as W is is positive, there is a systematic bias in equilibrium beliefs.. There are a number of intuitive results. The larger are Q, N and β, the more effectively the process eliminates distortions in the aggregation of information, the smaller is the variation in beliefs, and the more likely is consensus. In contrast, the larger is W, the larger is the variation in beliefs. W is, of course, the exogenous parameter that controls the degree to which the desire to minimize dissonance is achieved through altering networks as opposed to beliefs.

If Q, N and β work to promote the efficient social aggregation of information, W frustrates it.

From another perspective, ours is a model of the evolution of communication networks, which raises questions concerning the impact of exogenous variables in the model on the properties of equilibrium networks. We focus on two interrelated properties: the number of sects, and the overall connectivity of the network structure. Q has no discernible impact on network structure. The larger are N and β, the smaller is the number of sects, and the larger is the connectivity of the network structure. In contrast, the larger is W, the larger is the number of sects, and the less is the connectivity of the network structure. In other words, large W promotes a fragmented network structure.

Connecting the dots, we see two very different sorts of society: one that is fragmented and populated by people with conflicting beliefs which are often biased; another that is well connected and populated by people with convergent and unbiased beliefs. Our model suggests policies aimed at providing more information and a more connected society promote the later sort.

References

Acemoglu, D., Bimpikis, K. and Ozdaglar, A. (2010), Dynamics of information exchange in endogenous social networks, Technical report, National Bureau of Economic Research.

Acemoglu, D., Ozdaglar, A. and ParandehGheibi, A. (2010), 'Spread of (mis) information in social networks', *Games and Economic Behavior* **70**(2), 194–227.

Akerlof, G. A. and Dickens, W. T. (1982), 'The economic consequences of cognitive dissonance', *The American economic review* pp. 307–319.

Asch, S. E. (1955), 'Opinions and social pressure', *Scientific American* **193**, 35–35.

Bala, V. and Goyal, S. (1998), 'Learning from neighbors', *Review of Economic Study* **3**(65), 95–621.

Bala, V. and Goyal, S. (2000), 'A noncooperative model of network formation', *Econometrica* **68**(5), 1181–1229.

Berelson, B., Lazarsfeld, P. F., McPhee, W. N., Lazarsfeld, P. F. and Lazarsfeld, P. F. (1954), *Voting*.

Carsey, T. M. and Layman, G. C. (2006), 'Changing sides or changing minds? party identification and policy preferences in the american electorate', *American Journal of Political Science* **50**(2), 464–477.

DeGroot, M. H. (1974), 'Reaching a consensus', *Journal of the American Statistical Association* **69**(345), 118–121.

DeMarzo, P. M., Vayanos, D. and Zwiebel, J. (2003), 'Persuasion bias, social influence, and unidimensional opinions', *The Quarterly Journal of Economics* pp. 909–968.

Easley, D. and Kleinberg, J. (2010), *Networks, crowds, and markets: Reasoning about a highly connected world*, Cambridge University Press.

Elliot, A. J. and Devine, P. G. (1994), 'On the motivational nature of cognitive dissonance: Dissonance as psychological discomfort.', *Journal of personality and social psychology* **67**(3), 382.

Festinger, L. (1957), *A theory of cognitive dissonance*, Vol. 2, Stanford university press.

Festinger, L. and Carlsmith, J. M. (1959), 'Cognitive consequences of forced compliance.', *The Journal of Abnormal and Social Psychology* **58**(2), 203.

Golub, B. and Jackson, M. O. (2010), 'Naive learning in social networks and the wisdom of crowds', *American Economic Journal: Microeconomics* **2**(1), 112–149.

Golub, B. and Jackson, M. O. (2012), 'How hoomphily affects the speed of learning and best response dynamics', *Quartely Journal of Economics* **127**(3), 1287–1338.

Higgins, E. T. (1999), 'Saying is believing effects: When sharing reality about something biases knowledge and evaluations', *Shared cognition in organizations: The management of knowledge* pp. 33–49.

Holme, P. and Newman, M. E. (2006), 'Nonequilibrium phase transition in the coevolution of networks and opinions', *Physical Review E* **74**(5), 056108.

Huston, T. L. and Levinger, G. (1978), 'Interpersonal attraction and relationships', *Annual review of psychology* **29**(1), 115–156.

Jackson, M. O. and Rogers, B. W. (2007), 'Meeting strangers and friends of friends: How random are social networks?', *The American economic review* pp. 890–915.

Jackson, M. O. and Watts, A. (2002), 'The evolution of social and economic networks', *Journal of Economic Theory* **106**(2), 265–295.

Kandel, D. B. (1978), 'Homophily, selection, and socialization in adolescent friendships', *American journal of Sociology* pp. 427–436.

Kuran, T. (1997), *Private truths, public lies: The social consequences of preference falsification*, Harvard University Press.

Lazarsfeld, P. F., Merton, R. K. et al. (1954), 'Friendship as a social process: A substantive and methodological analysis', *Freedom and control in modern society* **18**(1), 18–66.

McPherson, M., Smith-Lovin, L. and Cook, J. M. (2001), 'Birds of a feather: Homophily in social networks', *Annual review of sociology* pp. 415–444.

Nickerson, R. S. (1998), 'Confirmation bias: A ubiquitous phenomenon in many guises.', *Review of general psychology* **2**(2), 175.

Pan, Z. (2012), 'Opinions and networks: how do they effect each other', *Computational Economics* **39**(2), 157–171.

Rabin, M. (1994), 'Cognitive dissonance and social change', *Journal of Economic Behavior & Organization* **23**(2), 177–194.

Schelling, T. C. (1978), *Micromotives and macrobehavior*, WW Norton & Company.

Walker, G. (2014), Cognitive dissonance, social learning, and overconfidence.

Walker, G. (2015), The Coevolution of Beliefs and Networks, PhD thesis, Simon Fraser University.

Figures

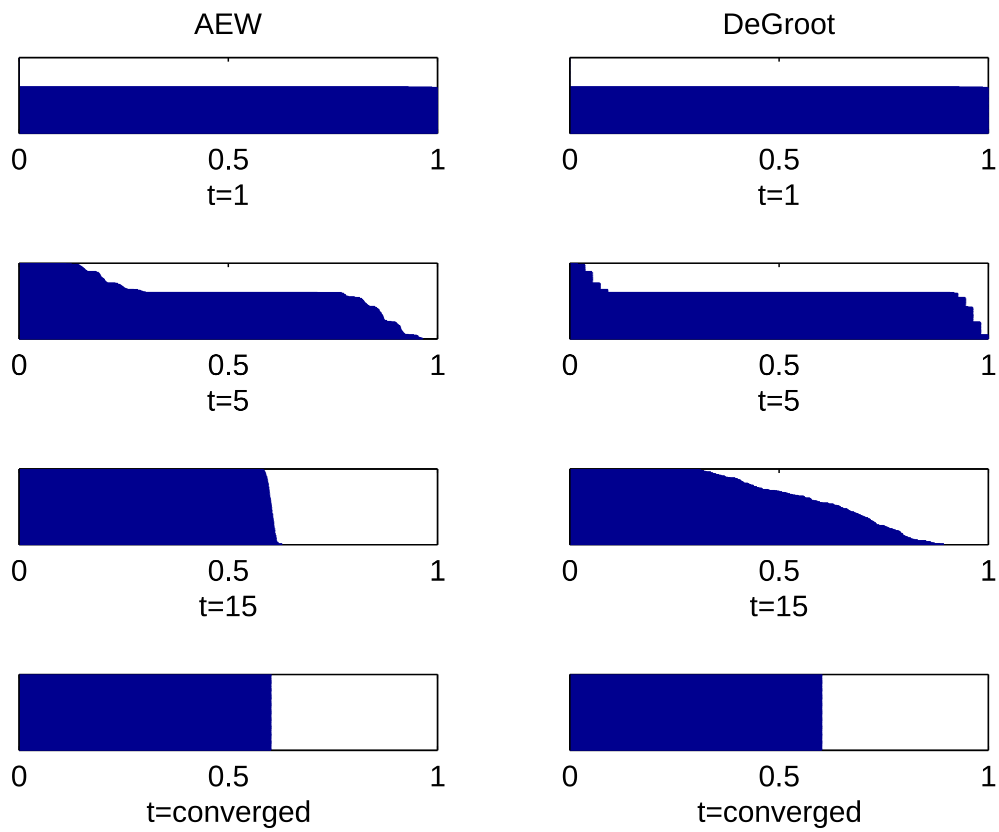

Figure 1: Evolution of Individual Beliefs without Network Changes

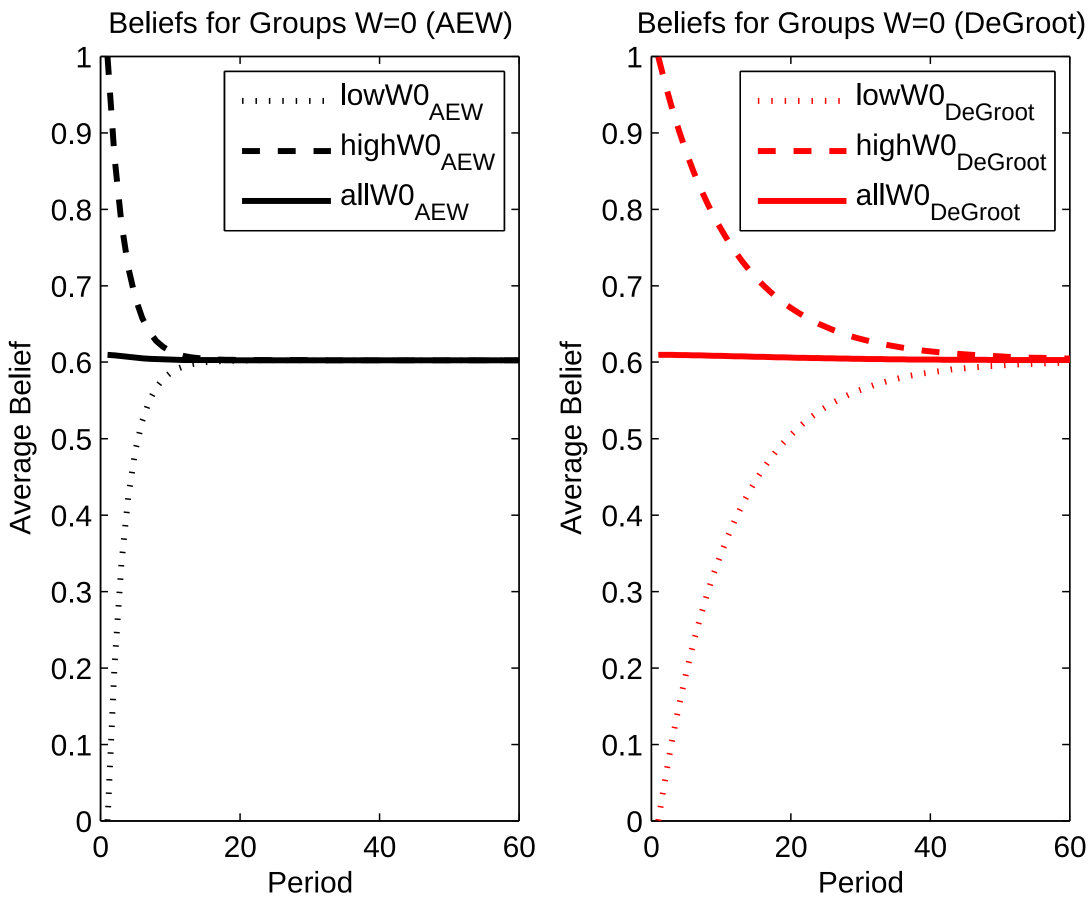

Figure 2: Evolution of Group Beliefs without Network Changes

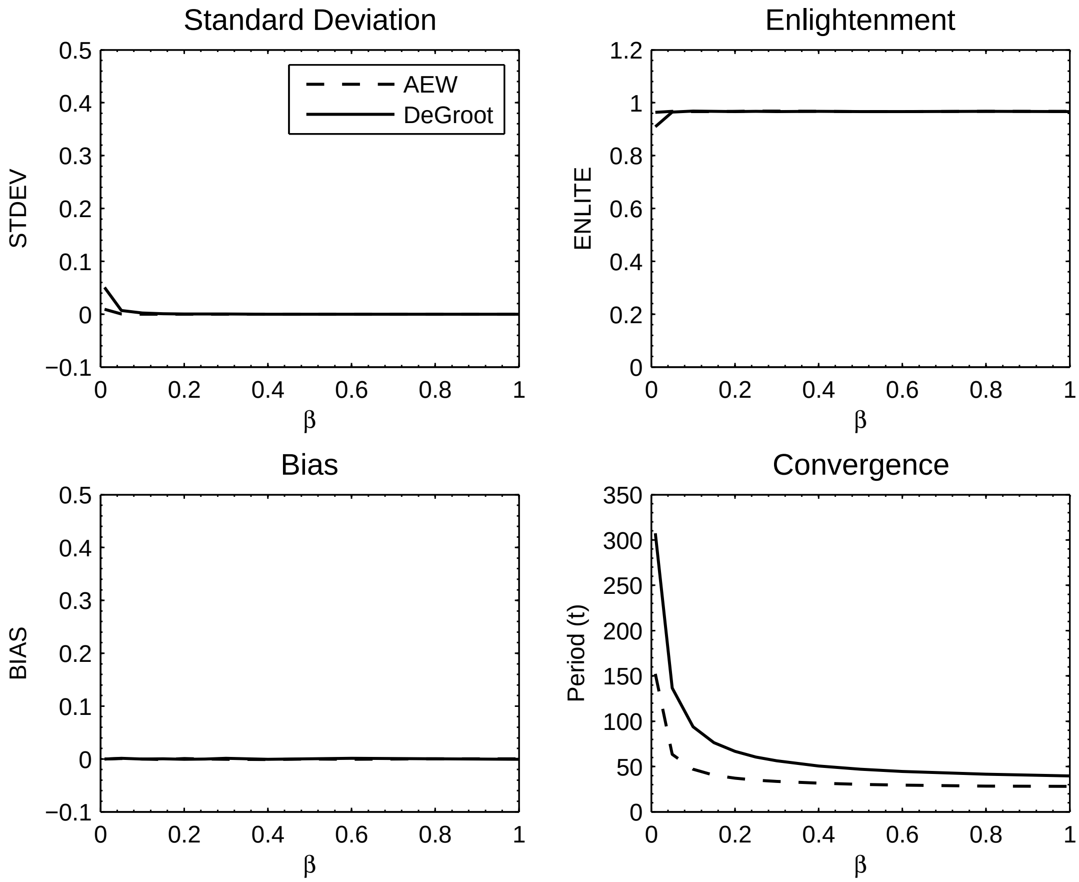

Figure 3: Belief-based Outcomes without Network Changes

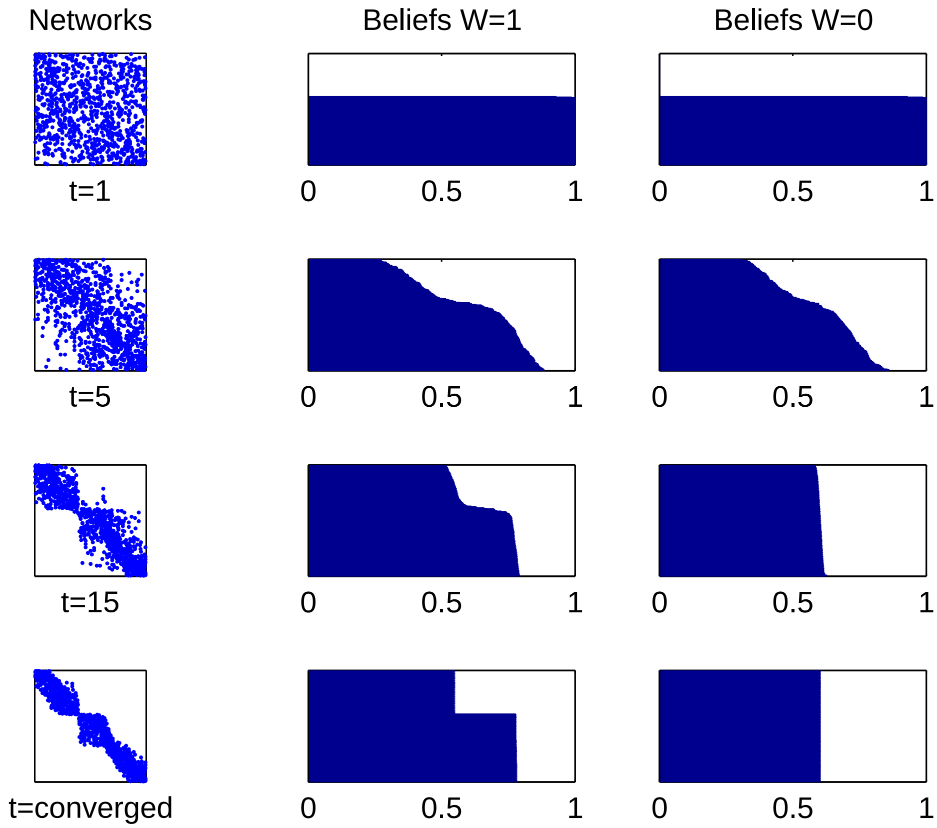

Figure 4: Coevolution of Beliefs and Networks

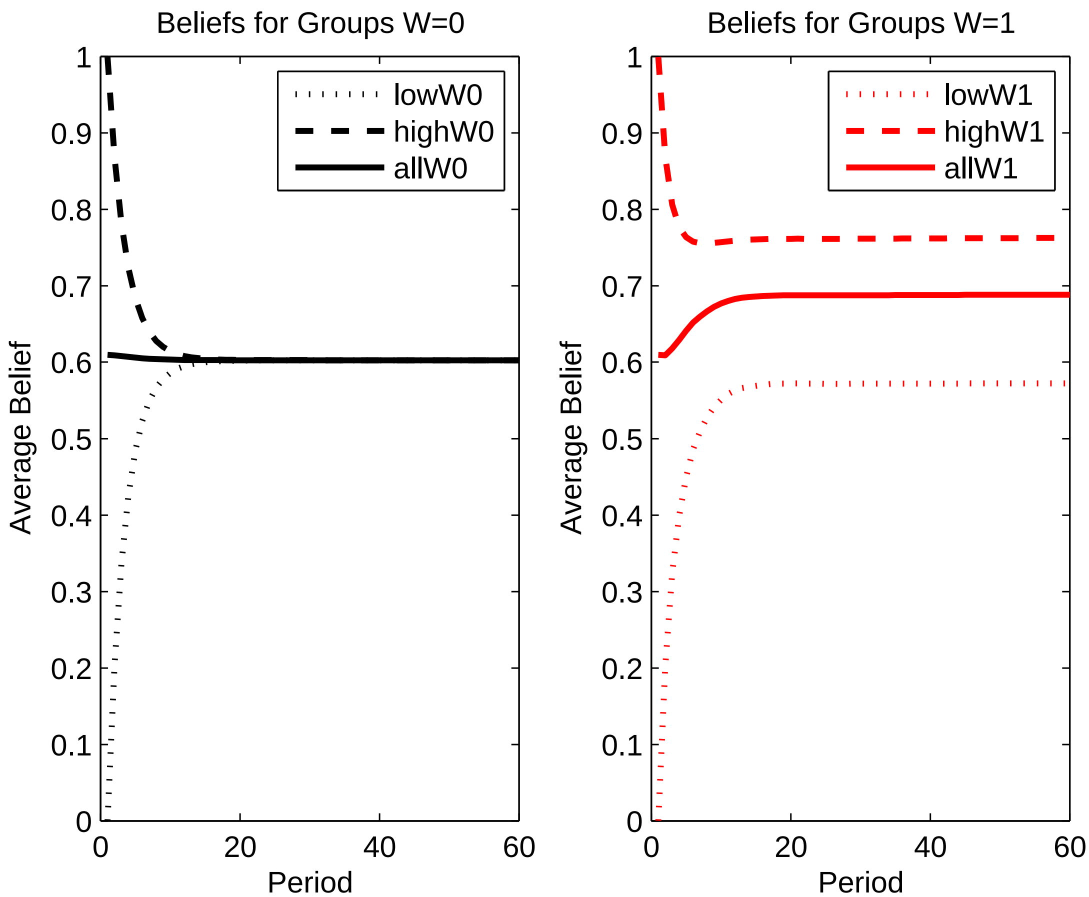

Figure 5: Evolution of Group Beliefs with Network Changes

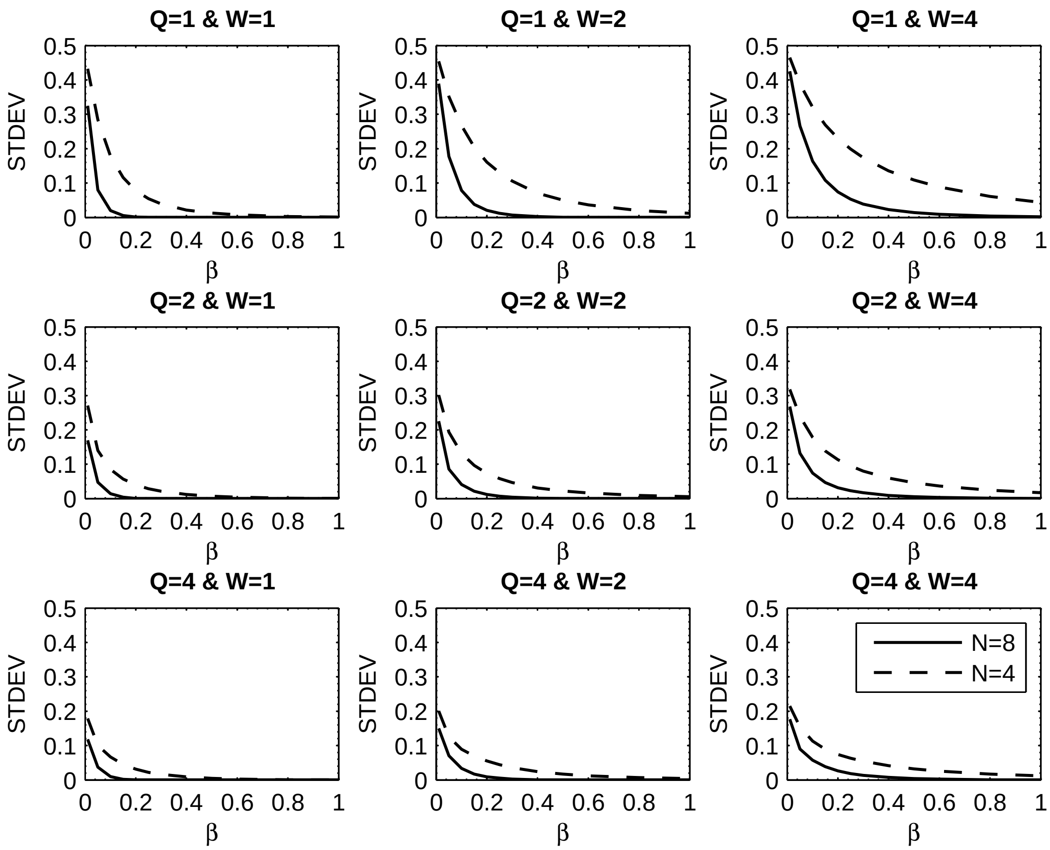

Figure 6: STDEV (Standard Deviation of Beliefs)

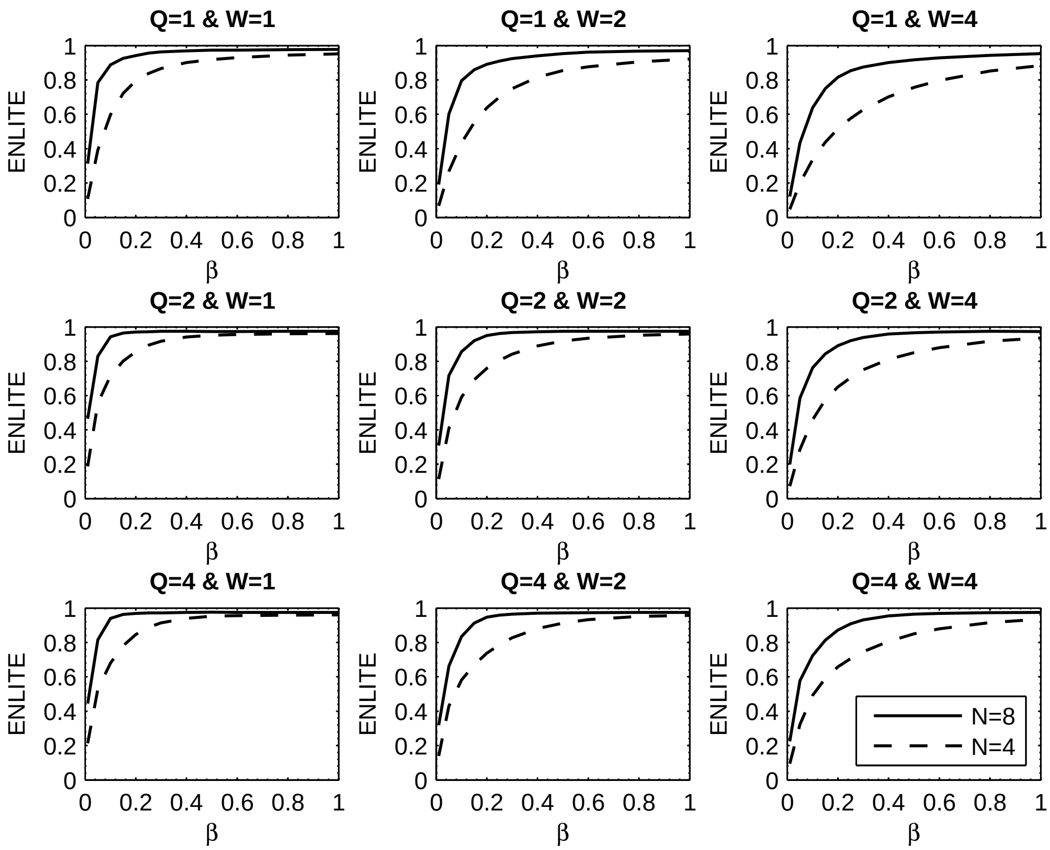

Figure 7: ENLITE (Enlightenment Index)

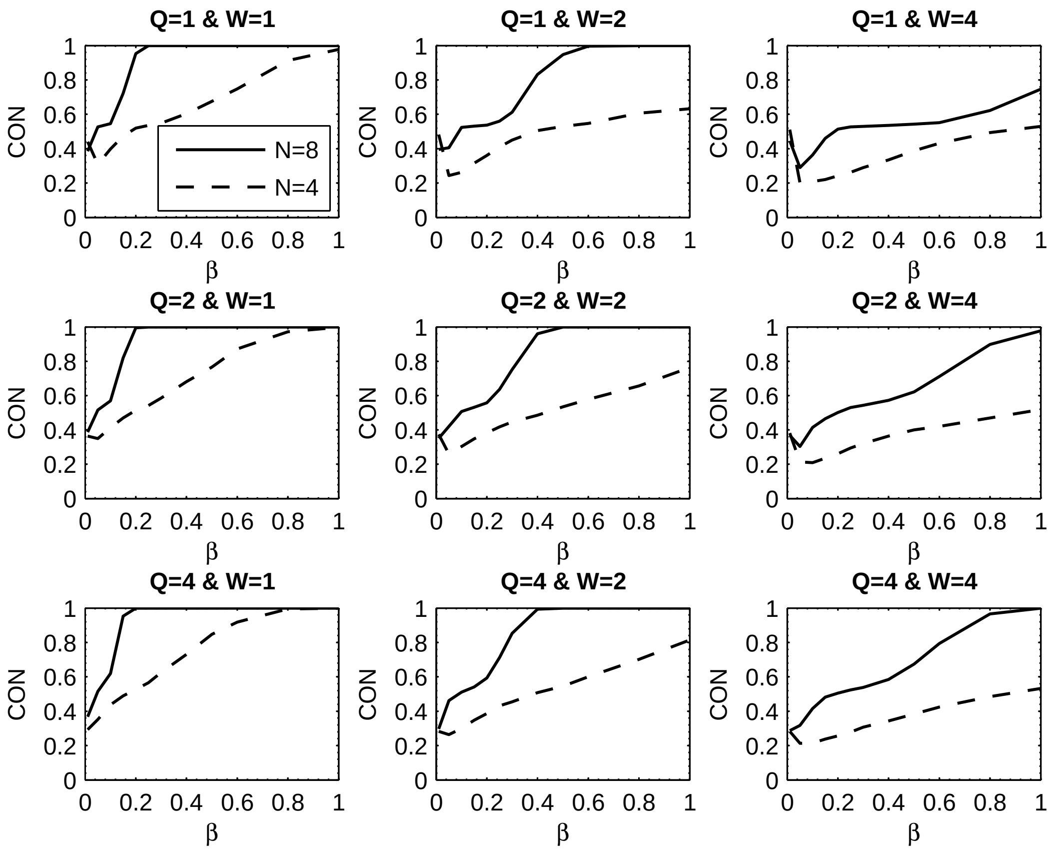

Figure 8: CON (Connectivity Index)

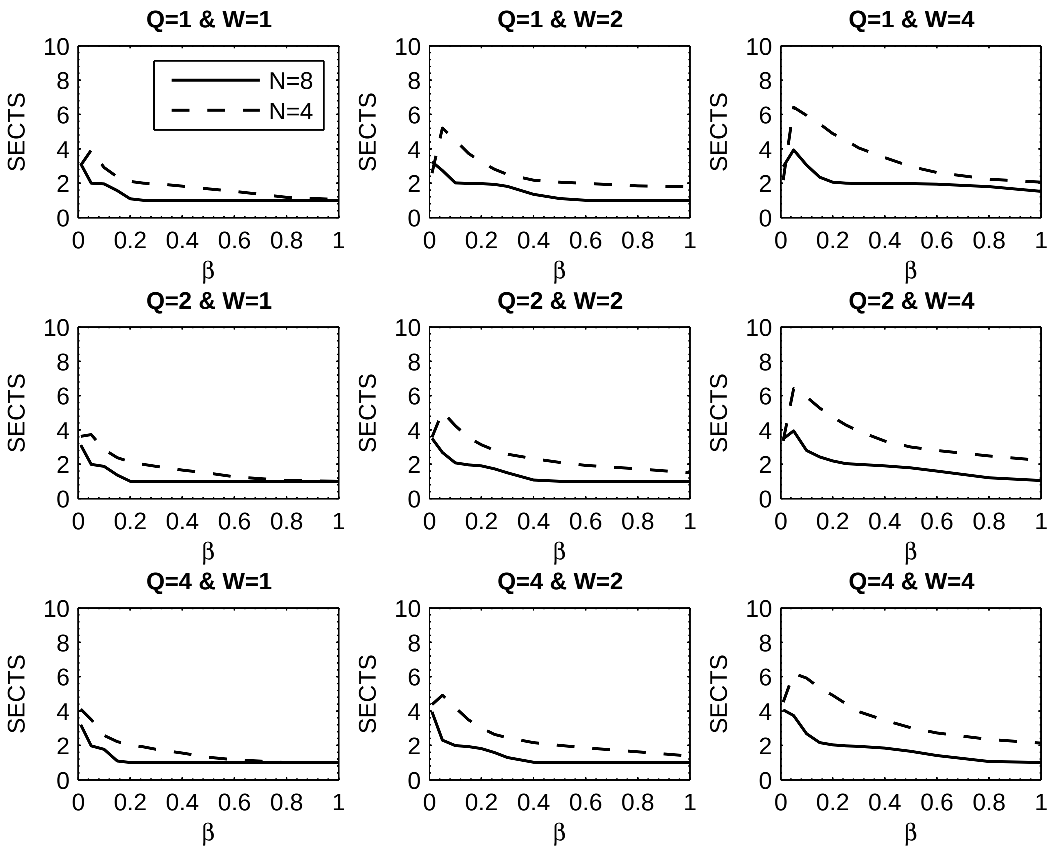

Figure 9: SECTS (Number of Sects)

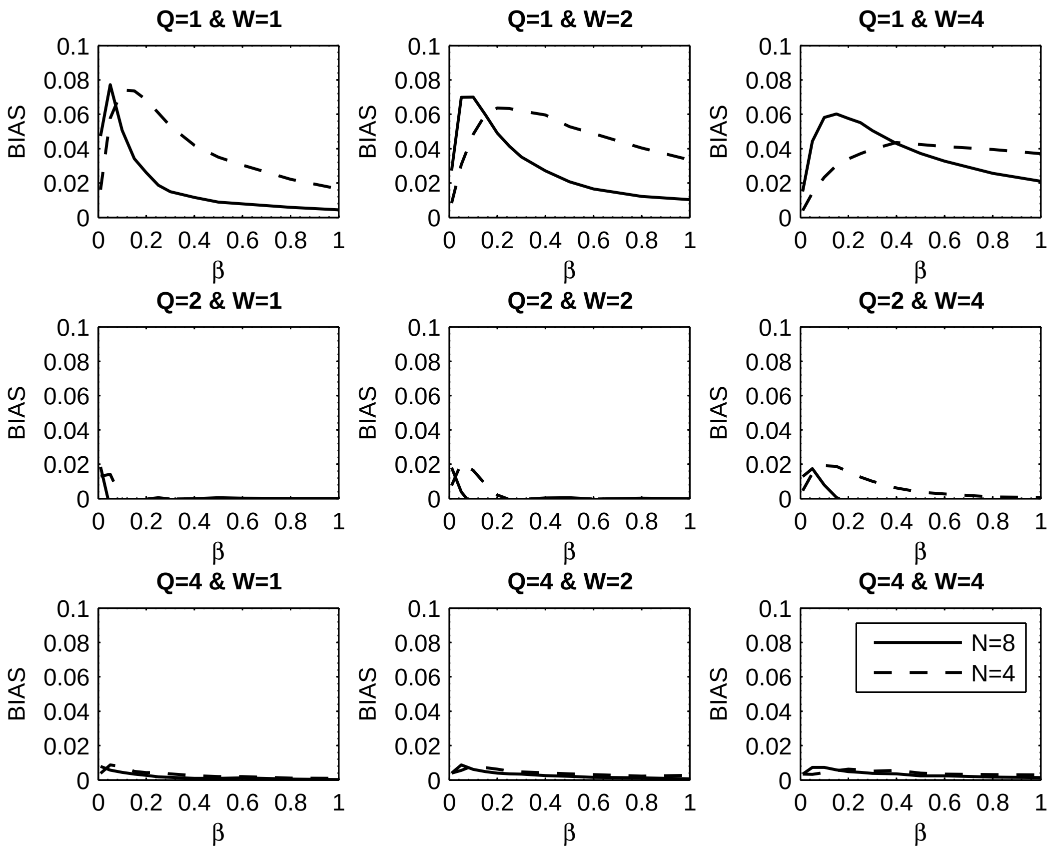

Figure 10: BIAS (Bias of Beliefs)

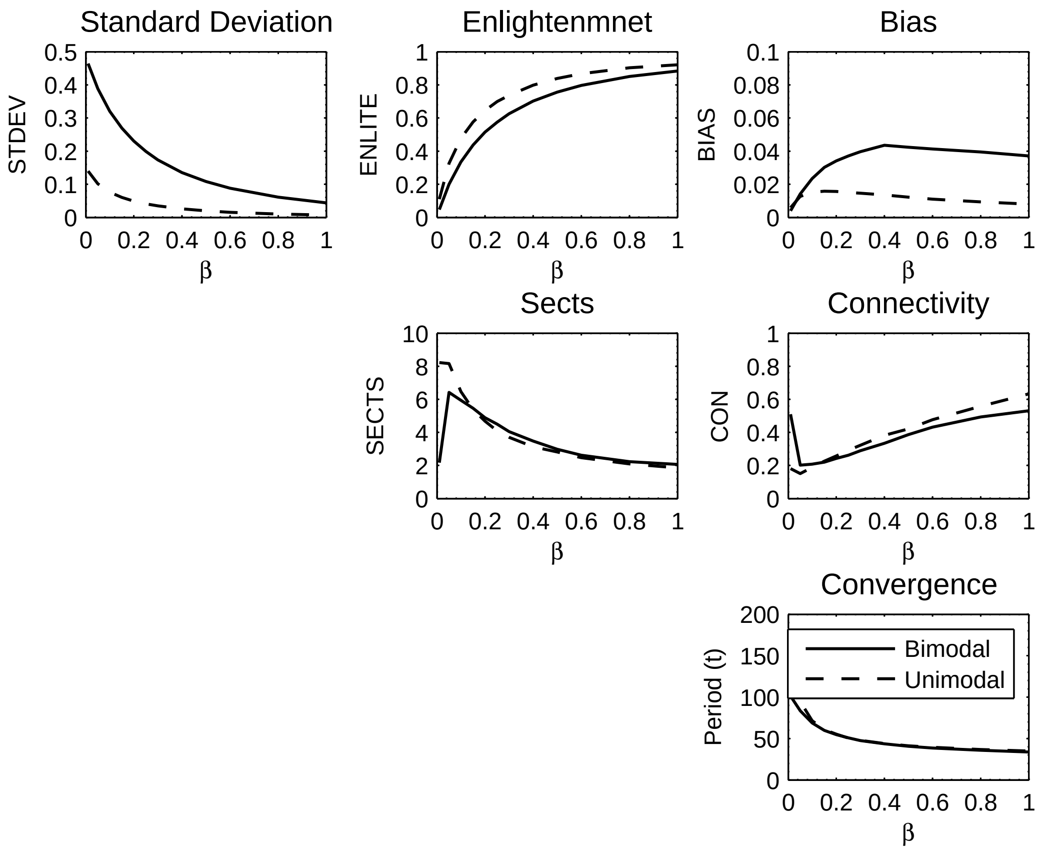

Figure 11: Bimodal vs Unimodal

Appendix: Tables with Means and 5th and 95th Percentiles

		N		4			8	
Q	β	W	1	2	4	1	2	4
			0.0491	0.1207	0.1980	0.0000	0.0092	0.0472
	0.2		0.0779	0.1611	0.2304	0.0012	0.0208	0.0742
			0.1061	0.1948	0.2568	0.0039	0.0323	0.0959
			0.0105	0.0422	0.1051	0.0000	0.0000	0.0099
1	0.4		0.0216	0.0703	0.1354	0.0000	0.0024	0.0227
			0.0332	0.0956	0.1613	0.0000	0.0055	0.0347
			0.0022	0.0168	0.0596	0.0000	0.0000	0.0030
	0.6		0.0076	0.0363	0.0882	0.0000	0.0001	0.0091
			0.0138	0.0544	0.1124	0.0000	0.0005	0.0157
			0.0228	0.0513	0.0928	0.0000	0.0065	0.0169
	0.2		0.0401	0.0736	0.1130	0.0003	0.0120	0.0316
			0.0578	0.0967	0.1328	0.0017	0.0176	0.0461
			0.0105	0.0165	0.0410	0.0000	0.0000	0.0034
2	0.4		0.0117	0.0309	0.0594	0.0000	0.0008	0.0091
			0.0193	0.0465	0.0765	0.0000	0.0028	0.0148
			0.0001	0.0072	0.0236	0.0000	0.0000	0.0000
	0.6		0.0076	0.0363	0.0882	0.0000	0.0000	0.0032
			0.0080	0.0254	0.0502	0.0000	0.0000	0.0063
			0.0212	0.0426	0.0619	0.0000	0.0056	0.0193
	0.2		0.0313	0.0552	0.0739	0.0000	0.0090	0.0263
			0.0422	0.0675	0.0873	0.0002	0.0125	0.0335
			0.0041	0.0161	0.0315	0.0000	0.0000	0.0044
4	0.4		0.0216	0.0703	0.1354	0.0000	0.0024	0.0227
			0.0140	0.0338	0.0524	0.0000	0.0015	0.0106
			0.0000	0.0067	0.0167	0.0000	0.0000	0.0000
	0.6		0.0026	0.0121	0.0253	0.0000	0.0000	0.0023
			0.0055	0.0177	0.0346	0.0000	0.0000	0.0044

Table 1: STDEV (Standard Deviation of Beliefs)

		N	4			8		
Q	β	W	1	2	4	1	2	4
1	0.2		0.7419	0.5935	0.4765	0.8846	0.8190	0.7887
			0.7957	0.6374	0.5145	0.9422	0.8904	0.8154
			0.8385	0.6815	0.5550	0.9919	0.9424	0.8395
	0.4		0.8186	0.7668	0.6628	0.9269	0.8872	0.8431
			0.9004	0.8159	0.7018	0.9686	0.9405	0.9012
			0.9553	0.8522	0.7403	0.9977	0.9901	0.9400
	0.6		0.8552	0.8007	0.7651	0.9370	0.9135	0.8748
			0.9302	0.8766	0.7975	0.9738	0.9616	0.9282
			0.9834	0.9219	0.8300	0.9982	0.9966	0.9721
2	0.2		0.8023	0.6879	0.5963	0.9259	0.9157	0.8484
			0.8558	0.7593	0.6512	0.9702	0.9495	0.8906
			0.9099	0.8291	0.7107	0.9963	0.9711	0.9328
	0.4		0.8186	0.8430	0.7581	0.9382	0.9338	0.9246
			0.9418	0.8895	0.8113	0.9743	0.9725	0.9583
			0.9753	0.9329	0.8683	0.9977	0.9968	0.9810
	0.6		0.8957	0.8952	0.8342	0.9360	0.9353	0.9303
			0.9302	0.8766	0.7975	0.9734	0.9746	0.9705
			0.9899	0.9663	0.9212	0.9973	0.9984	0.9914
4	0.2		0.7968	0.6777	0.6075	0.9259	0.9086	0.8403
			0.8460	0.7383	0.6588	0.9692	0.9466	0.8715
			0.8954	0.8029	0.7102	0.9978	0.9682	0.9038
	0.4		0.8893	0.8356	0.7550	0.9379	0.9299	0.9156
			0.9004	0.8159	0.7018	0.9686	0.9405	0.9012
			0.9712	0.9213	0.8584	0.9981	0.9971	0.9756
	0.6		0.9000	0.8943	0.8364	0.9377	0.9352	0.9284
			0.9561	0.9328	0.8788	0.9748	0.9729	0.9695
			0.9894	0.9634	0.9200	0.9984	0.9974	0.9923

Table 2: ENLITE (Enlightenment Index)

		N	4			8		
Q	β	W	1	2	4	1	2	4
1	0.2		0.3404	0.2559	0.1712	0.5007	0.4999	0.3482
			0.5190	0.3612	0.2424	0.9528	0.5373	0.5142
			0.5992	0.5397	0.3420	1.0000	0.5739	0.5590
	0.4		0.4977	0.3378	0.2203	1.0000	0.4993	0.5016
			0.6040	0.5056	0.3339	1.0000	0.8312	0.5348
			1.0000	0.5904	0.5166	1.0000	1.0000	0.5739
	0.6		0.4977	0.4975	0.2778	1.0000	1.0000	0.5007
			0.7483	0.5475	0.4318	1.0000	0.9960	0.5518
			1.0000	0.6608	0.5590	1.0000	1.0000	1.0000
2	0.2		0.3436	0.2563	0.1766	1.0000	0.4975	0.3372
			0.5142	0.3821	0.2616	0.9950	0.5579	0.5008
			0.6551	0.5608	0.4008	1.0000	1.0000	0.5820
	0.4		0.4977	0.3341	0.2298	1.0000	0.5007	0.4983
			0.6815	0.4860	0.3640	1.0000	0.9601	0.5720
			1.0000	0.6181	0.5522	1.0000	1.0000	1.0000
	0.6		0.4979	0.3731	0.2678	1.0000	1.0000	0.4979
			0.7483	0.5475	0.4318	1.0000	1.0000	0.7117
			1.0000	1.0000	0.5799	1.0000	1.0000	1.0000
4	0.2		0.3603	0.2540	0.1691	1.0000	0.4975	0.3973
			0.5261	0.3869	0.2562	1.0000	0.5932	0.5049
			1.0000	0.5120	0.3817	1.0000	1.0000	0.5289
	0.4		0.4977	0.3393	0.2266	1.0000	1.0000	0.4975
			0.6040	0.5056	0.3339	1.0000	0.8312	0.5348
			1.0000	1.0000	0.5047	1.0000	1.0000	1.0000
	0.6		0.4993	0.4975	0.2651	1.0000	1.0000	0.4977
			0.9184	0.6007	0.4245	1.0000	1.0000	0.7940
			1.0000	1.0000	0.5701	1.0000	1.0000	1.0000

Table 3: CON (Connectivity Index)

Q	β	N	4			8		
		W	1	2	4	1	2	4
			2.0000	2.0000	3.0000	1.0000	2.0000	2.0000
	0.2		2.0980	3.2300	4.8890	1.0950	1.9740	2.0550
			3.0000	4.0000	6.0000	2.0000	2.0000	3.0000
			1.0000	2.0000	2.0000	1.0000	1.0000	2.0000
1	0.4		1.8320	2.1780	3.4830	1.0000	1.3450	1.9830
			2.0000	3.0000	5.0000	1.0000	2.0000	2.0000
			1.0000	2.0000	2.0000	1.0000	1.0000	1.0000
	0.6		1.5190	1.9910	2.6170	1.0000	1.0080	1.9470
			2.0000	2.0000	4.0000	1.0000	1.0000	2.0000
			2.0000	2.0000	3.0000	1.0000	1.0000	2.0000
	0.2		2.1320	3.1350	4.7540	1.0100	1.8980	2.1980
			3.0000	4.0000	6.0000	1.0000	2.0000	3.0000
			1.0000	2.0000	2.0000	1.0000	1.0000	1.0000
2	0.4		1.6590	2.3160	3.3530	1.0000	1.0800	1.9050
			2.0000	3.0000	5.0000	1.0000	2.0000	2.0000
			1.0000	1.0000	2.0000	1.0000	1.0000	1.0000
	0.6		1.5190	1.9910	2.6170	1.0000	1.0000	1.5980
			2.0000	3.0000	4.0000	1.0000	1.0000	2.0000
			1.0000	2.0000	3.0000	1.0000	1.0000	2.0000
	0.2		2.0270	3.0100	4.9150	1.0000	1.8190	2.0360
			3.0000	4.0000	7.0000	1.0000	2.0000	3.0000
			1.0000	1.0000	2.0000	1.0000	1.0000	1.0000
4	0.4		1.8320	2.1780	3.4830	1.0000	1.3450	1.9830
			2.0000	3.0000	5.0000	1.0000	1.0000	2.0000
			1.0000	1.0000	2.0000	1.0000	1.0000	1.0000
	0.6		1.1650	1.8530	2.7310	1.0000	1.0000	1.4160
			2.0000	2.0000	4.0000	1.0000	1.0000	2.0000

Table 4: SECTS (Number of Sects)

Q	β	N	4			8		
		W	1	2	4	1	2	4
			0.0166	0.0128	0.0031	-0.0032	0.0086	0.0167
	0.2		0.0684	0.0636	0.0341	0.0261	0.0490	0.0576
			0.1163	0.1093	0.0619	0.0533	0.0835	0.0907
			-0.0062	0.0142	0.0096	-0.0119	-0.0006	0.0098
1	0.4		0.0419	0.0597	0.0435	0.0116	0.0271	0.0430
			0.0839	0.1022	0.0778	0.0344	0.0525	0.0724
			-0.0064	0.0049	0.0044	-0.0146	-0.0069	0.0038
	0.6		0.0305	0.0488	0.0413	0.0079	0.0165	0.0327
			0.0672	0.0917	0.0771	0.0302	0.0401	0.0591
			-0.0321	-0.0214	-0.0015	-0.0179	-0.0206	-0.0201
	0.2		-0.0058	0.0021	0.0157	-0.0002	-0.0014	-0.0022
			0.0231	0.0251	0.0350	0.0173	0.0177	0.0165
			-0.0062	-0.0273	-0.0141	-0.0151	-0.0153	-0.0191
2	0.4		-0.0035	-0.0033	0.0062	-0.0000	0.0004	-0.0027
			0.0242	0.0218	0.0280	0.0153	0.0161	0.0145
			-0.0274	-0.0263	-0.0198	-0.0156	-0.0157	-0.0183
	0.6		0.0305	0.0488	0.0413	0.0002	-0.0002	-0.0019
			0.0211	0.0228	0.0251	0.0155	0.0153	0.0138
			-0.0165	-0.0113	-0.0061	-0.0091	-0.0099	-0.0093
	0.2		0.0043	0.0063	0.0063	0.0027	0.0039	0.0048
			0.0252	0.0238	0.0185	0.0150	0.0185	0.0188
			-0.0170	-0.0140	-0.0093	-0.0096	-0.0086	-0.0086
4	0.4		0.0419	0.0597	0.0435	0.0116	0.0271	0.0430
			0.0213	0.0220	0.0200	0.0114	0.0141	0.0171
			-0.0170	-0.0140	-0.0123	-0.0094	-0.0098	-0.0090
	0.6		0.0019	0.0030	0.0034	0.0012	0.0015	0.0024
			0.0185	0.0201	0.0189	0.0123	0.0124	0.0144

Table 5: BIAS (Bias of Beliefs)