State-Dependent Opinion Dynamics*

Daron Acemoglu[1], Mohamed Mostagir[2], and Asuman Ozdaglar[2]

Abstract— **We study the simultaneous evolution of the opinion profile and network topology of a system of N agents. Based on the opinion profile at any given time, agents probabilistically decide which other agents to form links with. The probability of a link being formed with another agent depends on both similarity of their opinions and the popularity of that agent. Agents then average their opinion with the opinions of the agents they have formed links with, giving rise to a new opinion profile that determines –in a probabilistic fashion– the network topology for the next time step. Thus both opinions and network structure exhibit a strong correlation over time. Despite this correlation, we show that this system converges to a consensus in opinion. We provide a bound on the time to reach convergence, and simulate the limiting opinion profile as a function of the parameters of the system.**

I. INTRODUCTION

Individuals form beliefs on various economic, political and social variables based on information they receive from their local network. A growing literature studies dynamics of these beliefs using plausible rules of thumb on how individuals combine the beliefs of their neighbors in the network with their own—e.g., by taking some weighted average—and a fixed structure of the underlying social network (for example, Golub and Jackson in [1], DeGroot's seimnal work in [2]). This literature provides necessary and sufficient conditions for a consensus of opinion to emerge within a network. It is not clear, however, whether belief dynamics can be decoupled from the evolution of the network. For example, conformity of opinions between different agents may affect the likelihood that they will communicate in the future, and the popularity of individual might increase the weight that others attach to their opinion or whether they befriend them.

In this paper, we develop a tractable but fairly general model of the co-evolution of the topology of a social network and the distribution of beliefs over the network. In our model, agents form and sever links based on the similarity of beliefs and the past popularity of others in the network. These two features enable a fairly general process for the co-evolution of the topology of the network and the belief distribution.

Both of these features have been studied in isolation and in somewhat different contexts in the existing literature. Hegselmann-Krause introduced a model of belief dynamics in which agents only communicate with others in the network who are no further than a certain distance from their opinion [3]. This model has been subsequently studied in [4] and [5], among others, and variants where the impact of belief differences on communication probabilities is smooth have been analyzed in [6]. Skyrm and Pemantle, on the other hand, proposed a friendship model where past interactions between agents reinforce the chance of future interactions [7]. In [8], Fazeli and Jadbabaie studied an opinion formation process for a special case of the Skyrm-Pemantle model where interactions are governed by an urn process and a visit increases the probability of a future repeat visit in the same agents. Fazeli and Jadbabaie exploit the properties of the urn process to show convergence to a consensus.

Our work differs from these existing papers in several dimensions. First, we combine two dimensions of the endogeneity of the network, allowing future communications to depend both on differences in opinion and on the history of past visits. Second, we allow a tractable and general formulation linking these variables to communication probabilities (e.g., as opposed to Hegselmann-Krause's model where the impact of belief differences on these probabilities is discontinuous). As a result, compared to Fazeli and Jadbabaie, the probability that agent i visits agent j can increase without i having ever visited j. This may happen, for example, because j has visited another agent that holds an opinion close to that of i and as a result j's opinion has moved closer to that of i, increasing the probability of their future interaction. This is a plausible and quite realistic scenario, and one that is not captured by the Fazeli-Jadbabaie or Skyrm-Pemantle models. Compared to Hegselmann-Krause, communication is not precluded between any two agents because probabilities are continuous functions of belief differences and because the likelihood of communication with an agent holding a very different belief can be still very high if she is popular. These differences also imply that the mathematical structure of the evolution of the network topology and beliefs is different in our model than in existing work. In particular, the evolution of the network topology cannot be described by a stochastic process with independent and identically distributed increments. This necessitates a different mathematical approach, which we develop in the paper.

More formally, in our model at each date agents start with a scalar belief, which they use to determine which others to visit. After visiting another agent, an agent updates his beliefs by simply averaging their pre-visit belief with the belief of the agent that he visited. The model we consider relies on two parameters, β and κ, which have direct interpretations in terms of the nature of the social network. In particular, in our model β measures how "open" a society is to opinions that are different from the norm within that society, with higher values of β indicating a more "conservative" society. In a conservative society, agents are unlikely to visit other agents

* This research is partially supported by a Draper grant, ARO grant W911NF-09-1-0556, and AFOSR grant FA9550-09-1-0420.

[1] Department of Economics, MIT daron@mit.edu

[2] Laboratory for Information and Decision Systems, MIT {mosta,asu}@mit.edu

whose opinions are different from their own. The second parameter, κ, describes the tendency of a society to follow the opinions of popular individuals. As we show, these two parameters have a strong effect on the behavior of the system.

Because at time $t+1$ agents use opinions at the end of time t to determine what connections they will form at time $t+1$, the system is correlated across time. We show that despite this correlation and despite the fact that parameter values that can make it highly unlikely for agents with different opinions to interact, the system still reaches consensus, though the time to do so strongly varies as a function of the system parameters. We also derive bounds on the rate of convergence of beliefs to this consensus value. Finally, we provide simulations illustrating how the consensus value, which is a random convex combination of initial opinions, depends on the parameters β and κ.

The paper is organized as follows. Section II introduces our model formally and discusses the criteria by which agents decide to connect to other agents. Section III discusses differences between our model and existing consensus results and provides a convergence proof for opinion consensus. Section IV discusses how the system parameters affect the behavior of the system. Section V provides results from simulations, and Section VI concludes the paper.

II. Model

Our model consists of a society of N agents with an initial opinion vector $x(0) \in [0,1]^n$. Society is represented by a dynamic graph $G_t(N, E_t)$, where N is a fixed set of nodes representing the agents and E_t is the set of edges linking these agents at time t. The edges are formed by the agents in the following way. Each time step, all agents place weights on all other agents and use these weights to probabilistically select an agent to form a link with. We will refer to the link formation process as a 'visit'. Thus in each period each agent visits exactly one other agent and has an out-degree equal to one. We assume that the decision of whom to visit is independent across the agents. After each round, agents update their opinions. Agent i's opinion after a period where he visits agent j is the average opinion of i and j, while j's opinion remains unaffected by the visit. The weight that agents place on other agents has two components. The distance in opinion component, and the popularity component.

A. Distance in Opinion

The first component in the weight that agent i places on agent j at time t depends on how similar the opinions of i and j were at the end of the previous period, $t-1$, and is given explicitly by

$$w_{ij}(t) = e^{-\beta|x_i(t-1)-x_j(t-1)|}$$

where $\beta \geq 0$ is a constant. This expression is commonly used in the economics and physics literatures and lends itself naturally to our formulation, since agents with similar opinions put higher weights on each other. The probability with which agent i visits agent j at time t is then given by

$$\frac{e^{-\beta|x_i(t-1)-x_j(t-1)|}}{\sum_{k\neq i} e^{-\beta|x_i(t-1)-x_k(t-1)|}}$$

The parameter β can be agent specific, but for simplicity we take it to be the same for all agents. This parameter can be thought of as a quality that the society as a whole possesses. When β is high, agents are more inclined to visit only those agents who are similar to them in opinion, leading to a society that is more conservative on average than a society with lower values of β. When β is low, agents are more open to visit other agents who are different from them in opinion and society is more accepting of ideas that are different from the norm.

B. Popularity

Another aspect that goes into the weight that agents place on each other in our model is popularity. A popular agent at time t is an agent who has received many visits in time $t-1$. This means that popularity is equivalent to the in-degree of an agent, with agents with higher in-degrees being more popular. In particular, the weight that agent i assigns to agent j at time t based on popularity alone is equal to

$$w_{ij}(t) = e^{-\kappa(n-1-d_j(t-1))}$$

and is increasing in $d_j(t-1)$, the in-degree of agent j. Like before, the probability with which agent i visits j then is given by

$$\frac{e^{-\kappa(n-1-d_j(t-1))}}{\sum_{k\neq i} e^{-\kappa(n-1-d_k(t-1))}}$$

.

Similar to β, $\kappa \geq 0$ is a parameter that describes how important popularity is to agent when they assign weights to other agents. Again, this parameter can be agent specific but we take it to be the same across the society. This means that a society with high values of κ is inclined to visit popular agents more than a society with a lower κ.

Taking these two components together, we let $P(t)$ be the $n \times n$ probability matrix at time t, where entry $p_{ij}(t)$ represents the probability with which agent i visits agent j at time t, so that the i^{th} row represents the probability vector that agent i uses to decide whom to visit in period t. The probability $p_{ij}(t)$ is the combination of the two components described above. To reduce clutter, we drop the $t-1$ argument with the understanding that all weights are computed from the values obtained at the end of the previous period. Thus $P(t)$ is a matrix with $p_{ii}(t) = 0$ and

$$p_{ij}(t) = \frac{w_{ij}(t)}{\sum_{k\neq i} w_{ik}(t)} = \frac{e^{-\beta|x_i-x_j|-\kappa(n-d_j)}}{\sum_{k\neq i} e^{-\beta|x_i-x_k|-\kappa(n-d_k)}}. \quad (1)$$

At each time t, agents use $P(t)$ to decide whom to visit. The visits give rise to the set of edges E_t and the adjacency matrix $A(t)$. Because each agent visits exactly one other agent, the matrix $A(t)$ has one entry equal to one in each row with the rest of the entries in that row equal to zero. Since agents update their opinions by taking the average of their opinion as well as the opinion of the agent they visited, we define the stochastic matrix $Z(t)$ as

$$Z(t) = \frac{A(t) + I_n}{2},$$

where I_n is the $n \times n$ Identity matrix. We use the matrix $Z(t)$ to describe the opinion profile at time $t+1$ as a function of the opinion profile at time t through the relationship

$$x(t+1) = Z(t)x(t), \tag{2}$$

and therefore the opinion profile at time $t+1$ can be written as

$$Z(t)Z(t-1)...Z(0)x(0). \tag{3}$$

Let $\mathcal{Z} = \{Z : Z \in R^{n\times n}, Z_{ii} = \frac{1}{2}$ and $Z_{ij(i)} = \frac{1}{2}$ and $Z_{ik} = 0$ for $k \neq i, j(i)\}$ be the set of all such matrices.

III. CONSENSUS FORMATION

In this section we study the dynamics described by Eq. (2) and the backward product of the stochastic matrices $Z(t), ..., Z(0)$ in Eq. (3). We show that the system in Eq. (2) reaches consensus despite the strong correlation exhibited by the matrices $Z(t)$. A general result on the stability (and consensus) of systems like the one in Eq. (2) is given by Lorenz in [5]. The model we consider in this paper however violates at least one of the three conditions that Lorenz posits on the matrices $Z(t)$. Namely, the 'confidence is mutual' condition which states that the matrices $Z(t)$ should be type-symmetric. Type-symmetric matrices require that $Z_{ji}(t) > 0$ whenever $Z_{ij}(t) > 0$, which translates in our model to insisting that at time t, if agent i visits agent j then agent j must visit agent i. This is a restrictive assumption that our model does away with. This is not very surprising considering that Lorenz's conditions are sufficient but not necessary.

On the other hand, the infinite flow and absolute infinite flow properties developed in the work of Touri and Nedić on backward products of stochastic matrices provide necessary but not sufficient conditions for the backward product of the sequence $\{Z(t)\}_{t=0}^{\infty}$ to converge to a consensus matrix, i.e. checking whether our matrices fulfill these properties or not does not automatically lead us to conclusions about convergence. The characterization in [9] gives an equivalence between the absolute infinite flow property and the convergence of doubly-stochastic matrices. Since our matrices are not doubly-stochastic, we can not make use of this equivalence for our results. Finally, other convergence results starting from Wolfwitz's theorem [10], and found in the works of Jadbabaie, Lin, and Morse [11], and Hedrickx and Blondel [12] are based on the assumption that the matrices $Z(t)$ are irreducible, which is again not the case in our model.

As mentioned earlier, the closest model to ours is that of Fazeli and Jadbabaie. The proof of convergence in their work utilizes the fact that the model is built on top of a Polya's urn process, where the probability that agent i visits agent j evolves in the same way as in the urn process.[1] This enables the authors to use the fact that the urn process has the property that sequences are exchangeable, meaning that the order in which the visits are made does not affect the joint probability of a sequence of visits. Our model is more complicated because $p_{ij}(t)$ is not just a function of $p_{ij}(t-1)$, but also of who agent j visited in period $t-1$, since that visit will change x_j and therefore has an effect on $w_{ij}(t)$ and consequently $p_{ij}(t)$.

The following observation is important for our proof.

Observation 3.1: At any time step t and any $Z \in \mathcal{Z}$, We have $Pr(Z(t) = Z) > 0$.

This follows from how we define the weights and the probability matrix $P(t)$ in Eq. (1). Because $w_{ij}(t)$ is positive for all i, j, and t, the probability that any particular visit is made at time t is strictly positive. Since the visits are independent across all agents, the probability that any particular instance of visits, described by a matrix $A(t)$, is also strictly positive, and $Z(t)$ is always of the same sign as $A(t)$.

Before stating our convergence theorem, the following lemma, which builds on Observation 3.1, will be useful.

Lemma 3.2: Let $\mathcal{E}$ be a set of outcomes and $\{E_t\}_{t=0}^{\infty}$ be a sequence of events, not necessarily independent, such that $E_t \in \mathcal{E}$ and $Pr(E_t = E \in \mathcal{E}) > l$ for all E and t and some lowerbound $l > 0$. If $\sum_{t=0}^{\infty} Pr(E_t) = \infty$ then $Pr(\lim_{t\to\infty} \sup E_t) = 1$.

Proof: The proof relies on the following selection-rejection scheme, repeated at each time step t. At time t, let the probabilities of outcomes $E_1, E_2, ... \in \mathcal{E}$ be given by $p_1, p_2, ...$ and so on. Let $\mathcal{E}'$ be a nonempty subset of $\mathcal{E}$. Consider a sequence $\{W_t\}_{t=0}^{\infty}$ of random variables constructed as follows.

Step 1: With probability l, select an outcome E from $\mathcal{E}$. If $E \in \mathcal{E}'$, then set $W_t = E$ and stop. Else, go to step 2.

Step 2: Sample from $\mathcal{E}$ with probabilities $q_1, q_2, ...$ corresponding to $E_1, E_2, ... \in \mathcal{E}$, where $q_i = \frac{p_i - l}{1-l}$ for $E_i \in \mathcal{E}'$ and $q_i = \frac{p_i}{1-l}$ for $E_i \in \mathcal{E} \setminus \mathcal{E}'$. Set W_t equal to the resulting outcome.

What this scheme gives us is a sequence $\{W_t\}_{t=0}^{\infty}$ that mirrors the original sequence $\{E_t\}_{t=0}^{\infty}$ at each step in the sense that the probability of selecting any particular outcome is the same. However, selecting an outcome in $\mathcal{E}'$ with probability at least l is independent from previous events. Therefore we can use the Second Borel-Cantelli lemma since $\sum_{t=0}^{\infty} Pr(W_t | W_t \in \mathcal{E}') = \infty$ and the events $\{W_t | W_t \in \mathcal{E}'\}_{t=0}^{\infty}$ are independent, leading to $Pr(\lim_{t\to\infty} \sup W_t | W_t \in \mathcal{E}') = 1$. Since $\mathcal{E}'$ is arbitrary, the result follows. ■

The Second Borell-Cantelli lemma will be useful in the proof of our convergence theorem, but it cannot be used in its standard formulation since it assumes independence of the events E_t, an assumption that the matrices Z_t violate. What Lemma 3.2 does is provide an analogue for that lemma

[1] In this model, agent i visiting agent j increases the probability that this visit is repeated in the next period, in the same way that choosing a red ball from an urn and then putting back two red balls increases the probability of picking a red ball in the next trial.

when the events E_t are not independent but their probability of occurrence is always bounded away from zero.

For the purpose of stating the convergence theorem, we define a consensus matrix K as a matrix whose rows are all equal. We can then give the following statement.

Theorem 3.3: Denote by $Z(0,t)$ the matrix resulting from the multiplication $Z(t)Z(t-1)...Z(0)$ and let K be a consensus matrix, then $\lim_{t\to\infty} Z(0,t) = K$ with probability 1.

Proof: Following Hajnal in [13], we define the coefficient of ergodicity λ as

$$\lambda(Z(t)) = 1 - \min_{i,j} \sum_{k=1}^{n} \min\left(Z_{ik}(t), Z_{jk}(t)\right).$$

For a matrix Z, let us denote the second term above by

$$\delta_Z = \min_{i,j} \sum_{k=1}^{n} \min\left(Z_{ik}, Z_{jk}\right) \tag{4}$$

and the set of matrices

$$\mathcal{Z}_{\delta>0} \subset \mathcal{Z} \equiv \{Z : Z \in \mathcal{Z} \ and \ \delta_Z > 0\}.$$

Define the event

$$E_t : Z(t) \in \mathcal{Z}_{\delta>0}.$$

By Observation 3.1, $Pr(E_t) > 0$. Using Lemma 3.2 we have

$$Pr(\lim_{t\to\infty} \sup E_t) = 1,$$

i.e. the event E_t occurs infinitely many times. This implies that

$$\sum_{t=0}^{\infty} \delta_{Z(t)} = \infty. \tag{5}$$

The remaining step now is to show that the product $Z(t)Z(t-1)...Z(0)$ goes to K. Define the diameter d of a stochastic matrix Z as $d(Z)$, where

$$d(Z) = \max_{j} \left(\max_{i,k} \left(Z_{ij} - Z_{kj}\right)\right).$$

For two stochastic matrices Z_1 and Z_2, the shrinking lemma gives

$$d(Z_2 Z_1) \leq \lambda(Z_2) d(Z_1),$$

and hence we can write

$$\begin{aligned} d(Z(0,t+1)) &\leq (1-\delta_{Z(t)})d(Z(0,t)) \\ &\leq e^{-\delta_{Z(t)}} d(Z(0,t)) \\ &\leq e^{-\sum_{\tau=0}^{t} \delta_{Z(\tau)}} d(Z(0)). \end{aligned}$$

We have shown in Eq. (5) that $\sum_{t=0}^{\infty} \delta_{Z(t)} = \infty$ and therefore the last term goes to zero with probability 1. It is easy to see that for a stochastic matrix A, $d(A) = 0$ if and only if all the rows of A are the same, i.e. if $A = K$. Hence for the stochastic matrix $Z(0,t)$,

$$\lim_{t\to\infty} d(Z(0,t)) = 0 \to \lim_{t\to\infty} Z(0,t) = K.$$

∎

The previous result shows convergence regardless of the values of β and κ and of the fact that the evolution of the network and the evolution of opinions are highly correlated. When examined through the lens of Observation 3.1, this result can be understood in terms of the recurring "No man is an island" idea present in several works on opinion formation (for example, see [14]). This essentially means that, either directly or indirectly, communication is bound to happen between agents on the social network and hence with updating schemes like the one we consider, opinions eventually converge to a consensus. We will discuss this in more detail in the next section. We close this section by noting that the consensus matrix K is strictly positive, implying that all opinions factor to one degree or another in the limiting opinion profile. This is given in the next result.

Proposition 3.4: $\lim_{t\to\infty} Z(0,t) > 0$.

Proof: First consider the fact that for any two non-negative matrices A and B with positive diagonals, if $A_{ij} > 0$ or $B_{ij} > 0$ then $(AB)_{ij} > 0$. Therefore, as t increases, the number of zero entries in $Z(0,t)$ (monotonically) increases. Assume that there is t^* such that for all $t > t^*$, $Z(0,t) \not> 0$. In particular, consider an entry $Z(0,t)_{ij} = 0$ for all $t > t^*$. Because of Observation 3.1, $\exists t = \tau > t^*$ such that $Z(\tau)_{ij} > 0$, and hence $Z(0,\tau)_{ij} > 0$, contradicting our assumption. ∎

IV. RATES OF CONVERGENCE

We have seen that convergence is independent of the values of β and κ. However, as one would expect, the rate of convergence depends on these two parameters. Here, we use our convergence proof to obtain bounds on the convergence rate and to understand how the network evolves towards consensus. Although the bounds obtained by following the proof are rather loose, they highlight the role of the two parameters in the opinion formation process.

We start by examining the expression for δ_Z in Eq. (4). From the expression and assuming $n \geq 4$, one can verify that the set $\mathcal{Z}_{\delta>0}$ is equivalent to the set of matrices that have a positive column. The proof of the theorem requires that these matrices occur many times for the system to converge. However, the system can still converge without the help of the matrices in $\mathcal{Z}_{\delta>0}$. The occurrence of these matrices insures that all agents in the society obtain information from the same sources. But because G is connected in expectation, there are other ways for this to happen. Agent i who rarely visits agent j can have access to that agent's opinion through interacting with another agent k that interacts with j. This is equivalent to having a product of a string of matrices $Z = Z_1 \times Z_2 \times ...$ such that $Z_i \notin \mathcal{Z}_{\delta>0}$ but $Z \in \mathcal{Z}_{\delta>0}$. Therefore in reality one would expect the convergence rate to be faster[2] than the analysis provided here.

Formally, given the structure of the matrices in $\mathcal{Z}$, we can check that in Eq. (4), δ_Z can take one of only two values: 0, or $\frac{1}{2}$. This in turn means that $\lambda(Z(t))$ takes one of two

[2]Convergence will still not be fast if β is large, since if agent i visits j but not k, it indicates that i and k have similar opinions and therefore k is unlikely to visit j as well.

values, corresponding to δ_Z, either 1, in which case there is no shrinkage, or $\frac{1}{2}$.

Let $l^* = \sum_k \bar{P}(Z_k)$, where $Z_k \in \mathcal{Z}_{\delta>0}$ and $\bar{P}(Z_k)$ is a lower bound on the probability that Z_k is realized. Let M be the set of popular agents at time t, defined as $\{i | d_i(t) \geq d_j(t) \forall j\}$. For an agent $m \in M$, we bound the probability that an arbitrary agent i visits m, p_{im} in Eq. (1) as follows. We lower bound the numerator to get $e^{-\beta} \leq e^{-\beta|x_i - x_k|)}$. For the denominator, we have

$$\begin{aligned} \sum_j e^{(-\beta|x_i-x_j|+\kappa d_j)} &\leq \sum_j e^{(\kappa d_j)} \\ &= |M| e^{(d_{i\in M})} \\ &+ (N-|M|) e^{\kappa(d_{i\in M}-1)} \end{aligned}$$

where $|M|$ is the cardinality of M and $d_{i\in M}$ is the degree of popular agents. The last term follows from the fact that agents who are not in M must be at least one degree less than those who are. Simplifying the numerator and denominator we get

$$p_{ik} \geq \frac{e^{-\beta}}{|M| + (N-|M|)e^{-\kappa}}$$

and since $|M|$ can be at most equal to $N-1$, this simplifies to

$$p_{ik} \geq \frac{e^{-\beta}}{N-1+e^{-\kappa}} = l_0$$

Thus for any m in M we have $P(Z_m) \geq l_0^{N-1}$, so $l^* \geq l_0^{N-1}$. Now, choose some $\epsilon < l^*$ and refer to the event $Z(t) = Z_k \in \mathcal{Z}_{\delta>0}$ as a coupling. From Hoeffding inequality, we know that with probability greater than $1 - e^{-2\epsilon^2 T}$, there are at least $(l^* - \epsilon)T$ couplings happening between times 1 and T. Each time a coupling occurs, $\lambda(Z(t)) = \frac{1}{2}$ and the error gets halved, so with probability greater than $1 - e^{-2\epsilon^2 T}$ the error is less than $\frac{1}{2}^{T(l^*-\epsilon)}$, i.e. convergence is exponential, but because l^* can be quite small, the system can still take too long to converge.

The preceding analysis shows that convergence slows down as β increases and becomes faster with κ, as intuition would suggest. β slows down convergence because l^* decreases and the probability that $Z(t) \in \mathcal{Z}_{\delta>0}$. This is a direct result of conservatism, which prevents agents from agreeing on a common agent or a set of common agents to visit. When κ is high, the possibility of one of the coupling matrices being realized increases, since a column with all positive entries is equivalent to a popular agent who is visited by many agents. The analysis above is loose. One would expect the effect of κ to be more amplified and pronounced than β since κ is usually multiplied by the degree of the node. This contributes a high value to the node's weight, which in turn facilitates an increase in its probability of being visited and consequently, of convergence.

V. NUMERICAL SIMULATIONS

In this section, we simulate the opinion dynamics process described in Section II. [3] Our goal is to obtain a better understanding of the impact of the parameters β and κ on the speed of convergence and to obtain some insight into the limiting distribution.

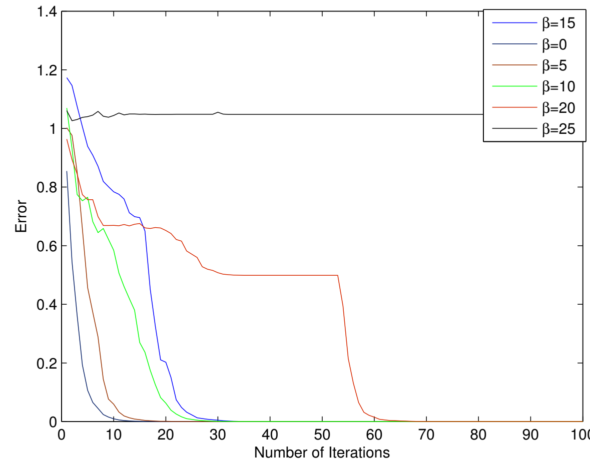

Fig. 1. Convergence Times For Different Values of β

A. Speed of Convergence

We simulate a system of $N = 20$ agents. Figure 1 shows the speed of convergence for different values of β. When $\beta = 0$, agents in the society have no particular inclination to visit those who are similar to them, the system converges exponentially fast, as one can verify from the previous section by setting $l^* = \frac{1}{n}$. As β grows larger convergence becomes slower. For β over 30, convergence was not obtained in any reasonable amount of time (1000000 iterations); instead, society clusters into two or three groups, with each group having a consensus amongst themselves.

Conversely, κ gives a strong boost to the speed of convergence. The way p_{ij} are formulated in our model implies that large values of κ can give a disproportionate weight to some agents, so much so that the value of the parameter β has no predictive value. This quickly leads the system to convergence as society obtains most of its information from only a small group of agents. Figure 1 is revisited in Figure 2 with the addition of $\kappa = 2$, which shows how the previously non-convergent case with $\beta = 20$ now converges extremely fast, as does the case when $\beta = 30$. At these values, the popularity component takes over and the system converges even when β is quite high.

B. Limit Opinion

We simulate the process for a randomly generated opinion vector with mean 0.351 and we start with the case $\beta = \kappa = 0$. This baseline case is when agents visit each other in each period with the same probability, thus the Z matrices are iid. This baseline case is analyzed in [15], where the authors show that the limiting opinion profile is a random vector with expectation equal to the expectation of the initial opinion profile, i.e. society is able to aggregate information (in this case, the average value) correctly. For a random opinion

[3] Code available on http://www.mit.edu/∼mostagir/

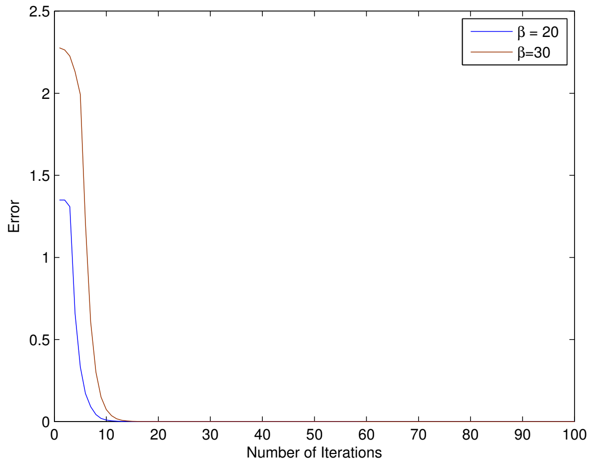

Fig. 2. Convergence Times for $\kappa = 2$ and $\beta = 20$ and 30

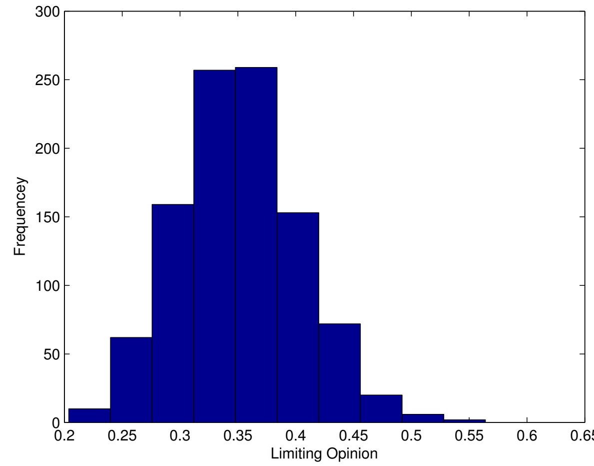

Fig. 3. Distribution of Expectation of Limit Opinion for $\beta = \kappa = 0$

profile, we ran the process 500 times, each run comprised 1000 time steps to insure convergence. The results are shown in Figure 3, where the mean of the resulting distribution approaches the expectation of $x(0)$.

When we keep κ equal to zero and increase β, the resulting expectation starts drifting away from the case $\beta = 0$. Figure 4 shows the expectation of the limit opinion falling steadily as β increases from 0 to 12. One explanation for this lies in the structure of the randomly generated opinion profile, which has five agents with $x(0)$ values that are less than 0.1, a 'majority' cluster at $t = 0$. This cluster can skew the resulting distribution by attracting the rest of the population closer to their position. Increasing β exacerbates this effect, leading to the conclusion that the initial distribution of opinions, not just in values but also in 'clusters' has an immediate effect on the limit distribution when β is high. A fringe majority cluster in the population can pull the rest of the network towards their opinion because they can practice their conservatism by visiting each other and ignoring the rest of the network, which given enough time, starts drifting towards the cluster. It is possible that this effect cancels out if there is a large number of agents, as groups of extreme agents on opposite sides neutralize each other, but the combination of a large number of agents and high β does not converge in a reasonable amount of time.

The behavior of the system as κ grows is unpredictable, as can be seen in Figure 5. For very high values like $\kappa = 10$ or more, the expectation of the limit can be anything that resembles the initial opinion of one of the agents in the population. A possible explanation for this is that an agent who gets just one more visit than the rest of the population early on gets a disproportionate amount of weight added to him since κ is too large, and dominates the process. When κ was set to lesser values like 0.2 and 0.5, 1, etc. the same unpredictable behavior was observed, in the sense that the expectation of the limit opinion did not display a meaningful correlation with κ. This is in contrast with the parameter β

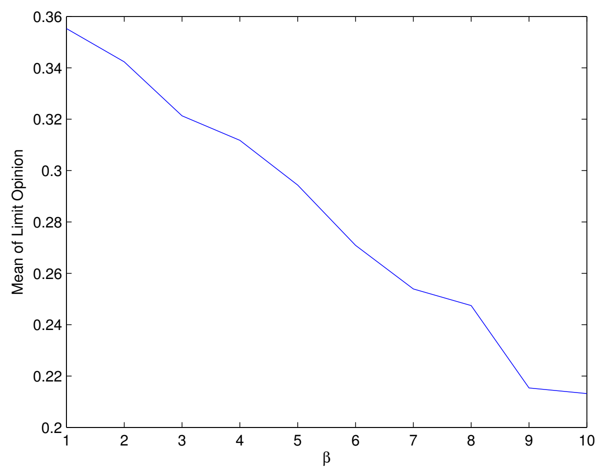

Fig. 4. Expectation of Limit Opinion as a Function of β

which either converged to a value that is close to the initial expectation when β was low, or systematically deviated from that value as β increased. The relationship between κ and the expectation of the limit opinion on the other hand is more volatile, making it difficult to make even a qualitative prediction about the resulting limiting opinion profile.

VI. Conclusion

In this paper, we presented a model of state-dependent opinion dynamics. Agents can form links with any other agents in the network, but are more likely to do so with those agents who carry a similar opinion to theirs. Agents also may want to connect to other agents based solely on their popularity in the society. We have shown that the resulting dynamical system converges to a consensus and that the rate of convergence correlates negatively with the level of conservatism in society and positively with popularity. Preliminary simulation results suggest that society loses its ability to combine information correctly as the level of

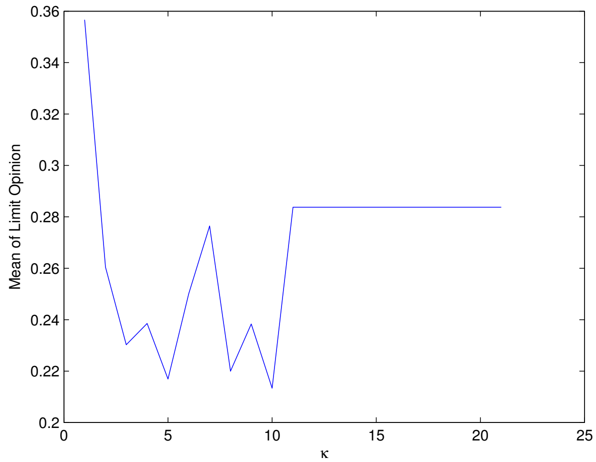

Fig. 5. Expectation of Limit Opinion as a Function of κ

conservatism increases, and that the so-called Super Star behavior, where agents are inclined to follow agents who are popular in society leads to uncertain outcomes. Current work focuses on deriving an analytical expression for the limiting opinion distribution. This is a difficult problem because the backward product of the matrices $Z(t)$ is not independent and does not have any exchangeability properties like those in [8], that can make it amenable to an iid analysis approach. We aim to develop techniques that enable us to approximate the limit distribution for the model we consider in this paper.

References

[1] B. Golub and M. Jackson, "Naive learning in social networks and the wisdom of crowds," *American Economic Journal: Microeconomics*, vol. 2, no. 1, pp. 112–149, 2010.

[2] M. DeGroot, "Reaching a consensus," *Journal of the American Statistical Association*, pp. 118–121, 1974.

[3] R. Hegselmann and U. Krause, "Opinion dynamics and bounded confidence: models, analysis and simulation," *Journal of Artificial Societies and Social Simulation*, vol. 5, no. 3, 2002.

[4] V. Blondel, J. Hendrickx, and J. Tsitsiklis, "On krause's multi-agent consensus model with state-dependent connectivity," *Automatic Control, IEEE Transactions on*, vol. 54, no. 11, pp. 2586–2597, 2009.

[5] J. Lorenz, "A stabilization theorem for dynamics of continuous opinions," *Physica A: Statistical Mechanics and its Applications*, vol. 355, no. 1, pp. 217–223, 2005.

[6] G. Como and F. Fagnani, "Scaling limits for continuous opinion dynamics systems." *Annals of applied probability: an official journal of the Institute of Mathematical Statistics*, vol. 21, no. 4, pp. 1537–1567, 2011.

[7] B. Skyrms and R. Pemantle, "A dynamic model of social network formation," *Proceedings of the National Academy of Sciences*, vol. 97, no. 16, p. 9340, 2000.

[8] A. Fazeli and A. Jadbabaie, "On consensus in a correlated model of network formation based on a polya urn process."

[9] B. Touri and A. Nedic, "Alternative characterization of ergodicity for doubly stochastic chains," in *IEEE Conference on Decision and Control*, 2011.

[10] J. Wolfowitz, "Products of indecomposable, aperiodic, stochastic matrices," *Proceedings of the American Mathematical Society*, vol. 15, pp. 733–736, 1963.

[11] A. Jadbabaie, J. Lin, and A. Morse, "Coordination of groups of mobile autonomous agents using nearest neighbor rules," *Automatic Control, IEEE Transactions on*, vol. 48, no. 6, pp. 988–1001, 2003.

[12] V. Blondel, J. Hendrickx, A. Olshevsky, and J. Tsitsiklis, "Convergence in multiagent coordination, consensus, and flocking," in *Decision and Control, 2005 and 2005 European Control Conference. CDC-ECC'05. 44th IEEE Conference on*. IEEE, 2005, pp. 2996–3000.

[13] J. Hajnal, "On products of non-negative matrices," in *Mathematical Proceedings of the Cambridge Philosophical Society*, vol. 79, no. 03. Cambridge Univ Press, 1976, pp. 521–530.

[14] D. Acemoglu, A. Ozdaglar, and A. ParandehGheibi, "Spread of (mis) information in social networks," *Games and Economic Behavior*, vol. 70, no. 2, pp. 194–227, 2010.

[15] A. Tahbaz-Salehi and A. Jadbabaie, "Consensus over ergodic stationary graph processes," *Automatic Control, IEEE Transactions on*, vol. 55, no. 1, pp. 225–230, 2010.