TEXT and rectangles in blue will NOT show on printed copy

Type final title of thesis or dissertation (M.S. and Ph.D.) below. If your title has changed since your submitted an Application for Graduate Degree, notify Graduate Office.

Agent-based Network Model of Social Impact with Bayesian Updating Rules

A Thesis

Presented to

the faculty of the School of Engineering and Applied Science

University of Virginia

in partial fulfillment

of the requirements for the degree

Master of Science

by

Name

Osama Eshera

Month degree is awarded May

Year 2015

APPROVAL SHEET

The thesis

is submitted in partial fulfillment of the requirements

for the degree of

Master of Science

AUTHOR signature

The thesis has been read and approved by the examining committee:

Please insert committee member names below:

Gerard P. Learmonth

Advisor

Alfredo Garcia

Paul Humphreys

Accepted for the School of Engineering and Applied Science:

Dean, School of Engineering and Applied Science

Month degree is awarded May

Year 2015

Print Form

Agent-based Network Model of Social Impact with Bayesian Updating Rules

---

A Thesis
Presented to
the faculty of the Sochol of Engineering and Applied Science
University of Virginia

---

in Partial Fulfillment
of the requirements for the degree
Master of Science

by

Osama M. Eshera

May 2015

Version: 1.7
Revised: April 15, 2015

Contents

1 Introduction

Complexity science is concerned with the characterization of natural phenomena through the postulation of generative rules which are unlike the laws of traditional science in that the latter aim to be universally predictive and are formulated through controlled experimentation and data analysis while the former aim to be locally descriptive and are often formulated through *in silico* simulation. These simulations, in turn, may display emergent properties which mimic natural phenomena. Complex systems are nonlinear, non-stationary, non-equilibrium, lack a characteristic scale (scale free), lack a central coordinating authority, and display some degree of order often on the verge of chaos.

The process of social influence and individual opinion formation can aptly be constructed as a complex adaptive system. Individual people formulate their opinions through some unknown process that considers social pressure and their own individual agency. What arises from this process is a global opinion distribution, or "public opinion," that is a nonlinear epiphenomena of local agent action. Opinion dynamics models postulate simple agent rules that approximate this unknown local process so as to generate global behavior mimicking observed phenomena.

1.1 Motivation

Developing a complex systems model of public opinion will have a twofold impact. First, it will provide a framework for a novel conceptualization of public opinion. Present understandings of public opinion rely on the linear aggregation of individual opinions, e.g., surveys from polling organizations such as Gallup or Pew Research Center. However, this fails to describe the process that drives such opinon formation in the first place. It neglects the influence of interest groups, political and media elites, and social mores, as well as individual variation on dimensions of knowledge, ideology, self-interest, group identity, emotion, and salient frames of reference, among others. A complex systems framework will shed light on the processes of social influence and individual response which will yield more fruitful results for questions about public opinion posed by political, non-profit, and community organizations.

1.2 Problem statement

Existing opinion dynamics models and theories are characterized by three significant faults. First, they largely assume that all people respond to social influence in the same manner. As a result, they use one-dimensional theories that are insufficient to explain individual variation and group dynamics. Second, in formulating said theories, existing approaches rely heavily on phenomena observed in biology or physics — such as spin interaction, bird flocking, or cellular automata — yet, there is little, if any, evidence that establishes a link between these quantitative models and human behavior. As a result, current models are often predetermined systems with intuitive results, given the initial conditions, and rarely display true emergent properties. Finally, the simulation spaces to which such models are applied are not representative of the true social spaces. Namely, most models employ a random graph topology despite there being strong evidence that human networks are not well approximated by random graphs.

1.3 Research aims

As a consequence of the aforementioned shortcomings, current opinion dynamics models offer little insight into human behavior. The broad objective of this research is to contribute to the development of a more descriptively rigorous and practically useful model of social influence. Towards this end, we have three research aims.

> *Aim 1:* Design a set of agent rules that is nuanced and accounts for variation in individuals' reception to and exercise of social influence. This rule set will take its roots in the study of social influence and individual response by psychologists and social scientists.
>
> *Aim 2:* Develop and employ a network topology that is a better approximation of human social networks.
>
> *Aim 3:* Determine issues (opinions) and phenomena whose dynamics can be fruitfully simulated using this approach.

2 Literature Review

For this project, two literature reviews are necessary: one surveying approaches to modeling opinion dynamics and a second outlining psychological theories of social influence. This section covers both bodies of literature.

2.1 Opinion dynamics models

The majority of opinion dynamics models belong to one of a few major categories: voter, averaging, majority rule, social impact, Sznajd, bounded confidence, and Bayesian updating models. The basic form and significant variations of each model type are presented here.

2.1.1 Voter model

The voter model was first introduced by Holley and Liggett (1975) as part of a study of the ergodic properties of a branching process. In each time-step of the opinion model, a node is randomly selected and then one of its neighbors (or itself) is also selected at random. Then, the selected node moves, invariably, to the opinion of its selected neighbor. In the default case, this model reaches consensus.

Several variations of the voter model have been studied, for instance, Yildiz et al. (2011) present a model with stubborn agents and show this prevents opinion convergence and invites further questions such as identifying optimal location of the stubborn agents. Xiong et al. (2010) consider a voter model in which agents own a memory in addition to a present opinion; they show the system settles to a more polarized state. Yang et al. (2012) posit a nonlinear voter model in which a tuning parameter controls the speed and nature of consensus formation. Finally, and perhaps most interestingly, Kempe et al. (2013) present an arbitrary graph structure in which agents are only open to influence from other agents with sufficiently similar interaction patterns and characterize various equilibria and end-states based on initial conditions.

2.1.2 Sznajd model

Sznajd-Weron (2005) held that changing a person's opinion is more likely to happen if that person is persuaded by two or more people than by a single individual. In its original version, agents occupy the sites of a linear chain, and have binary opinions, denoted by Ising spin variables. A pair of neighboring agents, i and $i+1$, determines the opinions of their two nearest neighbors, $i-1$ and $i+2$, as follows: if the agents of the pair share the same opinion, they successfully impose their opinion on their neighbors. If, instead, the two agents disagree, each agent imposes its opinion on the other agent's neighbor. Opinions are updated in a random sequential order. Starting from a totally random initial configuration, where both opinions are equally distributed, two types of stationary states are found, corresponding to consensus, with all spins up ($m = 1$) or all spins down ($m = -1$), and to a stalemate, with the same number of up and down spins in antiferromagnetic order ($m = 0$).

2.1.3 Averaging model

Often used to study the emergence of consensus in a society, agents in this model update their opinions based on the average of their neighbors' opinions as introduce by DeGroot (1974). Opinions in this model may be represented as either discrete or continuous. Friedkin and Johnsen (1990) extended this model to encompass both dissent and consensus by implementing a two-tiered opinion model in which each agent owns both a public and private opinion.

Variations of the averaging model include the study of exogenous inputs and defiant agents with local and global interactions by Fotouhi and Rabbat (2013) in which they model local interaction as agents' response to direct neighbors (family, friends, etc.) and global interaction as agents' response to indirect influence (mass media, advertising, etc.). Each node finds the proportion of local and global agents with whom it disagrees and probabilistically determines whether or not to adopt their average opinion. To model mass media, in particular, an exogenous bias on the average system orientation is added. Fotouhi and Rabbat present the results of both simulations then investigate the effect of stubborn nodes that refuse influence by the exogenous bias. Their key result is the necessary conditions so that the stubborn nodes can successfully resist influence by mass media.

2.1.4 Bounded confidence model

These models are structurally similar to the averaging models but differ in that agents are only influenced by neighbors whose opinions are within a given confidence range of their own opinion. Hegselmann and Krause (2002) find the system will approach consensus only if the network is such that for any two agents i and j there exists a third one k such that a chain of confidence leads from i to k as well as from j to k. Opinions in this model lie on the real line.

Numerous studies have examined the convergence of bounded confidence models to consensus or clustering. One such variation is the multi-equilibria regulation model in which Koulouris et al. (2012) explore the effect of various network topologies — complete graph, star, cellular automata, small-world, and random graph — on opinion dynamics in a bounded confidence model. The results of their simulations are quite nuanced and can be found summarized in Castellano et al. (2009).

2.1.5 Majority rule model

Majority models differ from averaging and bounded confidence models in that selected agents adopt the opinion of the majority of other agents (including themselves). In each time step, the focal agent probabilistically determines whether or not to adopt the opinion of the majority of his neighbors. Zhang et al. (2013) show that, in the default case, this will lead to consensus with two system sub-states: opinion consensus and relation consensus. They show the number of opinion clusters in the final state is inversely proportional to the number opinions in the initial state.

The non-consensus opinion (NCO) model was proposed by Shao et al. (2009) in which they found clusters of nodes holding the same opinion occurs when the concentration of nodes holding the same opinion is above a certain threshold. Shao et al. (2009) show these clusters are impermeable to influence from others. Li et al. (2012) generalize the NCO model by adding a weight factor to each individual's original opinion to control the relative displacement in opinion at the end state. Li et al. (2012) find tuning this parameter will impact the stability and opinion minority density of the clusters. They present various other detailed results in.

2.1.6 Continuous opinions and discrete actions

The continuous opinions and discrete actions (CODA) model was proposed by Martins (2007) and generalized as a Bayesian updating procedure in Martins (2012). This model differentiates between opinion and choice: agents possess a discrete action that is motivated by an underlying continuous opinion distribution in which he registers information received from neighbors and keeps track of his own memory. In this approach, agents with more extreme opinions will require more influence to change their public opinion. In Martins (2008), there are two possible discrete actions to choose from and a simple application of Bayes rule is used as the updating procedure for agents' underlying continuous opinion distribution.

In Martins (2007) this decision theory is outlined and applied to the discrete voter model and shows the discrete model can be obtained as a limit case of the CODA model. In Martins (2012) this decision theory is applied to the continuous bounded confidence model which is also shown to be a special case of the CODA model. Martins (2008) explores the behavior of extreme opinions in the CODA model and finds that extreme opinions do not survive long with increasing social contact. Liu et al. (2012) present this model with multiple cluster-coupled patterns in which clusters are connected via fixed links and interact with each other at varying frequencies. They show that higher interconnectivity and interaction frequency yield to greater levels of agreement and decreased extremism.

2.2 Psychological Theories of Social Influence

There is, of course, a vast body of literature in which psychologists and social scientists attempt to explain some aspect of individual opinion formation as a function of social influence. The majority of such studies employ a controlled experiment methodology and are, thus, only able to comment on specific aspects of an individual's opinion formation process: gender, minority status, religious affiliation, etc.

There are; however, a few studies that postulate a generalized theory of social influence, a few of which are summarized here.

2.2.1 Dissemination of culture model

This model was introduced by Axelrod (1997) and depends on two critical aspects of cultural dynamics: social assimilation and homophily. Social assimilation is the tendency of individuals to become more similar when they interact. Homophilly is the tendency of likes to attract each other, so that they interact more frequently. The remarkable result of this model is that, while social scientists expected these two aspects to yield a self-reinforcing dynamic leading to a global convergence, Axelrod and others demonstrated the persistence of local diversity.

In this model, agents are represented as nodes in a network and are given F integer variables $(\sigma_1, \ldots, \sigma_F)$ that can assume q possible values, i.e. $\sigma_f = 0, 1, \ldots, q-1$. The variables are "cultural features" and q is the number of the possible traits allowed per feature. Together, they represent the different "beliefs, attitudes, and behavior" of agents. In an elementary dynamic step, an individual i and one of his neighbors j are selected and interact with probability $\omega_{i,j}$ where $\delta_{i,j}$ is Kronecker's delta.

$$\omega_{i,j} = \frac{1}{F} \sum_{f=1}^{F} \delta_{\sigma_f(i), \sigma_f(j)} \tag{1}$$

If the interaction does, in fact, take place, one of the features for which i and j have different traits, i.e. $\sigma_f(i) \neq \sigma_f(j)$, is selected and $\sigma_f(j)$ is set equal to $\sigma_f(i)$. For any pair of neighbors, there are two stable configurations. First is the case in which they are exactly equal, so that they belong to the same cultural region. Second is the case in which they are entirely different, so that they exist at the fringe of cultural regions.

2.2.2 Social impact theory

The psychological theory of social impact was first introduced by Latane (1981) to characterize the process through which people respond to and exert influence on their peers. The impact of a social group on a subject depends on three factors: (1) the number of neighbors, (2) their convincing power, and (3) the distance from the subject, where the distance can be interpreted spatially or abstractly.

Each agent, i, has an opinion $\sigma_i = \pm 1$, a measure of persuasiveness p_i, and a measure of supportiveness s_i. Persuasiveness and supportiveness characterize the agent's ability to convince his neighbors to change or keep their opinion, respectively. p_i and s_i take random values. Between two agents exists a distance d_{ij}. The total impact that an agent experiences from his environment is given by:

$$I_i = \left[\sum_{j=1}^{N} \frac{p_j}{d_{ij}^{\alpha}} (1 - \sigma_i \sigma_j) \right] - \left[\sum_{j=1}^{N} \frac{s_j}{d_{ij}^{\alpha}} (1 + \sigma_i \sigma_j) \right] \tag{2}$$

Where $\alpha > 2$ expresses how fast the impact decreases with the distance. The first term represents the persuasive impact while the second represents the supportive impact. An agent changes his opinion if the persuasive pressures exceed supportive pressures, or:

$$\sigma_i(t+1) = -\text{sgn}\left[\sigma_i(t)I_i(t) + h_i\right] \tag{3}$$

Where h_i is a random field representing all pressures other than social impact that may persuade or support the current opinion. This field is most often set to zero for simplicity.

3 Model theoretics and methodology

3.1 Agent rules

The rules employed to describe the low-level behavior of agents in the previously surveyed models vary widely depending on the objective of the simulation. The goal of this model is to bridge between statistical physics approaches and social science theories. A preliminary rule set will begin the social impact theory and then make a number of incremental additions and adaptations to enhance the model and explore it's suitability for mimicking various phenomena.

3.1.1 Social impact theory

As described above, the social impact theory assigns each person a measure of persuasion – their ability to convince those with opinions differing from their own to change their position and adopt one matching their own position. Likewise, each person is assigned a measure of supportiveness – their ability to convince those with opinions identical to their own to retain their shared position inspite of persuasive influence from others. Latane (1981) argues that these characteristics, while related, are independent. So, each node, i in the network is seperately assigned p_i and s_i from a random uniform distribution with values between 0 and 100.

Each node in the network also requies an initial opinion. In this implementation of the model, the possible opinions are $+1$ and -1. Each node is initially assigned one of the possible opinions with equal probability.

Between each pair of nodes, or set of neighbors, is some measure of distance. This is represented as a weight on the edges in the network. Absent any information about the nature of each node and the nature of its relationship to its neighbor, this must also be assigned arbitrarily. In this case, each edge from node i to node j carries a weight d_{ij} drawn from a random uniform distribution with values between 1 and 10. It is worth noting here, all edges in this network are undirected, implying that all relationships are mutual.

In this implementation of the social impact theory, the random field h_i is set to zero.

Finally, the value for α is set arbitrarily and will be the subject a more detailed discussion in section 3.3 on simulation methods.

The aforementioned parameters are set at the beginning of the simulation and do not change. The model was updated synchronously, in each iteration a new copy of the graph is created in which all changes are made based on the original graph. At the end of iteration, the new graph becomes the old graph and a new copy is generated. The algorithm passes over each node, computes I_i as given in Equation

(2) above, and then determines $\sigma_i(t+1)$ as given in equation(3) above. $\sigma_i(t+1)$ are stored in a new copy of the graph.

The first term of Equation (2) is the additive persuasive impact of the neighbors of node i while the second term is the additive supportive impact of the neighbors of node i – when nodes i and j share the same opinion, that is $\sigma_i = \sigma_j$, the persuasive impact goes to zero and the supportive impact persists. Likewise, when node i and j hold opposing positions, that is $\sigma_i \neq \sigma_j$, the supportive impact goes to zero and the persuasive impact persists. By Equation (1), if the overall persuasive impact exceeds the overall supportive impact, node i will change it's position.

3.1.2 Social impact theory with memory

In a simple, yet significant, enhancment of this model we consider agents with memory. The model follows the same form as the previously described approach except after computing the opinion dynamics rule in Equation (3) the result is not immediately adopted as the agent's new opinion. Instead, the result is appended to a unique list for each agent. When communicating a discrete opinion, each agent may report some point estimate from that distribution – the mean, the mode, a random draw, a random draw from the more salient or recent memories.

The objective with such a modification to the model is to dampen the longterm variability in agent opinions. That is, agents are less likely to make dramatic changes in their opinions over short periods of time and agents become less likely to change their opinions altogether in the longrun.

For the consideration of memory to really be fruitful, agents must have the capacity to hold and be exposed to more than just binary opinions.

3.1.3 Social impact theory with three opinion state space

In the standard model, opinions take value $\sigma_i = \pm 1$. This implies people may only take position for or against a particular question. In fact, people are more often ambivalent or apathetic than opinionated. In this adaptation of the social impact theory, $\sigma_i \in \{-1, 0, +1\}$ where 0 represents a neutral position.

To implement this change, we must first introduce a new node parameter alongside persuasion, p_i, and supportiveness, s_i: ambivalence, a_i. This is a measure of how grounded a person is in his/her ambivalence, given that they hold the neutral position. s_i was assigned from a uniform random distribution between 0 and 100.

Computing social impact, I_i now becomes more intricate. If $\sigma_i = \pm 1$ and if $\sigma_j = \pm 1$, the computation is identical to that detailed earlier. But, in the case when either $\sigma_i = 0$ or $\sigma_j = 0$, it is easiest to consider I_i in its three components: persuasive impact, $I_{p,i}(t)$, supportive impact $I_{s,i}(t)$, and apathetic impact, $I_{a,i}(t)$.

The process for computing $I_i(t)$ can be summarized as follows:

- In each iteration t, node i keeps track of persuasive impact, $I_{p,i}(t)$, supportive impact $I_{s,i}(t)$, and apathetic impact, $I_{a,i}(t)$ while surveying each of its neighbors.
- If $\sigma_i(t) = +1$:
  - If $\sigma_j(t) = +1$:
    - $I_{s,i}(t) = I_{s,i}(t) + \frac{s_j}{d_{ij}^\alpha}$ $\quad I_{p,i}(t) = I_{p,i}(t) + 0$ $\quad I_{a,i}(t) = I_{a,i}(t) + 0$
  - If $\sigma_j(t) = -1$:
    - $I_{s,i}(t) = I_{s,i}(t) + 0$ $\quad I_{p,i}(t) = I_{p,i}(t) + \frac{p_j}{d_{ij}^\alpha}$ $\quad I_{a,i}(t) = I_{a,i}(t) + 0$
  - If $\sigma_j(t) = 0$:
    - $I_{s,i}(t) = I_{s,i}(t) + 0$ $\quad I_{p,i}(t) = I_{p,i}(t) + 0$ $\quad I_{a,i}(t) = I_{a,i}(t) + \frac{a_j}{d_{ij}^\alpha}$
- If $\sigma_i(t) = -1$:
  - If $\sigma_j(t) = +1$:
    - $I_{s,i}(t) = I_{s,i}(t) + 0$ $\quad I_{p,i}(t) = I_{p,i}(t) + \frac{p_j}{d_{ij}^\alpha}$ $\quad I_{a,i}(t) = I_{a,i}(t) + 0$
  - If $\sigma_j(t) = -1$:
    - $I_{s,i}(t) = I_{s,i}(t) + \frac{s_j}{d_{ij}^\alpha}$ $\quad I_{p,i}(t) = I_{p,i}(t) + 0$ $\quad I_{a,i}(t) = I_{a,i}(t) + 0$
  - If $\sigma_j(t) = 0$:
    - $I_{s,i}(t) = I_{s,i}(t) + 0$ $\quad I_{p,i}(t) = I_{p,i}(t) + 0$ $\quad I_{a,i}(t) = I_{a,i}(t) + \frac{a_j}{d_{ij}^\alpha}$
- If $\sigma_i(t) = 0$:
  - If $\sigma_j(t) = +1$:
    - $I_{s,i}(t) = I_{s,i}(t) + 0$ $\quad I^+_{p,i}(t) = I^+_{p,i}(t) + \frac{p_j}{d_{ij}^\alpha}$ $\quad I_{a,i}(t) = I_{a,i}(t) + 0$
  - If $\sigma_j(t) = -1$:
    - $I_{s,i}(t) = I_{s,i}(t) + 0$ $\quad I^-_{p,i}(t) = I^-_{p,i}(t) + \frac{p_j}{d_{ij}^\alpha}$ $\quad I_{a,i}(t) = I_{a,i}(t) + 0$
  - If $\sigma_j(t) = 0$:
    - $I_{s,i}(t) = I_{s,i}(t) + 0$ $\quad I_{p,i}(t) = I_{p,i}(t) + 0$ $\quad I_{a,i}(t) = I_{a,i}(t) + \frac{a_j}{d_{ij}^\alpha}$

Here, $I^-_{p,i}(t)$ and $I^+_{p,i}(t)$ are the persuasive impacts from neighbors whose opinion is -1 and $+1$ respectively. The opinion dynamics decision will also differ from that which was previously presented. At the end of the iteration, $\sigma_i(t+1)$ is determined based on persuasive impact $I_{p,i}(t)$, supportive impact $I_{s,i}(t)$, and apathetic impact, $I_{a,i}(t)$ detailed above.

- If $\sigma_i(t) = +1$:
  - If $max\{I_{p,i}(t), I_{s,i}(t), I_{a,i}(t)\} = I_{s,i}(t)$: $\quad \sigma_i(t+1) = +1$
  - If $max\{I_{p,i}(t), I_{s,i}(t), I_{a,i}(t)\} = I_{p,i}(t)$: $\quad \sigma_i(t+1) = -1$
  - If $max\{I_{p,i}(t), I_{s,i}(t), I_{a,i}(t)\} = I_{a,i}(t)$: $\quad \sigma_i(t+1) = 0$

- If $\sigma_i(t) = -1$:
  - If $max\{I_{p,i}(t), I_{s,i}(t), I_{a,i}(t)\} = I_{s,i}(t)$: $\sigma_i(t+1) = -1$
  - If $max\{I_{p,i}(t), I_{s,i}(t), I_{a,i}(t)\} = I_{p,i}(t)$: $\sigma_i(t+1) = +1$
  - If $max\{I_{p,i}(t), I_{s,i}(t), I_{a,i}(t)\} = I_{a,i}(t)$: $\sigma_i(t+1) = 0$
- If $\sigma_i(t) = 0$:
  - If $max\{I^+_{p,i}(t), I^-_{p,i}(t), I_{a,i}(t)\} = I^+_{p,i}(t)$: $\sigma_i(t+1) = 0$
  - If $max\{I^+_{p,i}(t), I^-_{p,i}(t), I_{a,i}(t)\} = I^-_{p,i}(t)$: $\sigma_i(t+1) = -1$
  - If $max\{I^+_{p,i}(t), I^-_{p,i}(t), I_{a,i}(t)\} = I_{a,i}(t)$: $\sigma_i(t+1) = 0$

3.1.4 Social impact theory with Bayesian updating

Now we consider a further abstraction from the binary opinion space. It is still true that $\sigma \in \{-1, 0, +1\}$ but now agents do not simply have a discrete opinion, instead, they have a probability distribution, $f_i(\sigma)$ over all possible opinions. Initially, this opinion distribution is assigned randomly (but, as outlined below, this distribution will be updated in each iteration).

This opinion distribution is kept private, when communicating their opinion agents report the position, σ_i, which they feel is most likely to be correct – that is, the position with the highest probability in the distribution. The objective of each agent is to improve its inference about σ^*.

To do so, the agent updates its distribution in accordance to Bayes theorem. The focal agent, i, surveys each neighbor, j, one at a time and receives a communication σ_j. The focal agent must have in mind some relationship between its belief in the true value of σ, call it σ^*, and the value communicated by its neighbor, σ_j. This relationship is given by a probability distribution $Pr(\sigma_j|\sigma^*)$ which expresses the chance that neighbor j will communicate σ_j given that a possible value σ^* is the correct value (in the mind of the focal agent).

In the context of the social impact theory, $Pr(\sigma_j|\sigma^*)$ can be determined through an adaptation of the impact measure expressed in Equation (2). In this case, the measures of persuasion, supportiveness, and apathy must be conceived of differently than before.

$$\begin{aligned} Pr(\sigma_j = +1|\sigma^* = -1) &= p_j \\ Pr(\sigma_j = -1|\sigma^* = -1) &= s_j \\ Pr(\sigma_j = 0|\sigma^* = -1) &= 1 - p_j - s_j \end{aligned} \tag{4}$$

$$\begin{aligned} Pr(\sigma_j = +1|\sigma^* = +1) &= s_j \\ Pr(\sigma_j = -1|\sigma^* = +1) &= p_j \\ Pr(\sigma_j = 0|\sigma^* = +1) &= 1 - p_j - s_j \end{aligned} \tag{5}$$

$$\begin{aligned} Pr(\sigma_j = +1|\sigma^* = 0) &= p_j \\ Pr(\sigma_j = -1|\sigma^* = 0) &= p_j \\ Pr(\sigma_j = 0|\sigma^* = 0) &= 1 - p_j - p_j \end{aligned} \tag{6}$$

In each iteration of the algorithm, a focal node, i, is chosen at random and then i randomly selects one of its neighbors, j, and receives a communication from j, σ_j, indicating the position j believes is most likely to be true. The focal agent can then update its underlying distribution as follows:

$$f_i(\sigma^*|\sigma_j) = \frac{Pr(\sigma_j|\sigma^*)Pr(\sigma^*)}{Pr(\sigma_j)} \tag{7}$$

In the next iteration of the algorithm, this posterior distribution becomes the new prior distibution and the algorithm repeats. For the sake of completeness, a complete set of update expressions contained in Equation (7).

$$f_i(\sigma = +1|\sigma_j = +1) = \frac{Pr(\sigma_j = +1|\sigma = +1)Pr(\sigma = +1)}{\left[\begin{array}{c}Pr(\sigma = +1)Pr(\sigma_j = +1|\sigma = +1)+ \\ Pr(\sigma = -1)Pr(\sigma_j = +1|\sigma = -1)+ \\ Pr(\sigma = 0)Pr(\sigma_j = +1|\sigma = 0)\end{array}\right]} \tag{8}$$

$$f_i(\sigma = -1|\sigma_j = +1) = \frac{Pr(\sigma_j = +1|\sigma = -1)Pr(\sigma = -1)}{\left[\begin{array}{c}Pr(\sigma = +1)Pr(\sigma_j = +1|\sigma = +1)+ \\ Pr(\sigma = -1)Pr(\sigma_j = +1|\sigma = -1)+ \\ Pr(\sigma = 0)Pr(\sigma_j = +1|\sigma = 0)\end{array}\right]} \tag{9}$$

$$f_i(\sigma = 0|\sigma_j = +1) = \frac{Pr(\sigma_j = +1|\sigma = 0)Pr(\sigma = 0)}{\left[\begin{array}{c}Pr(\sigma = +1)Pr(\sigma_j = +1|\sigma = +1)+ \\ Pr(\sigma = -1)Pr(\sigma_j = +1|\sigma = -1)+ \\ Pr(\sigma = 0)Pr(\sigma_j = +1|\sigma = 0)\end{array}\right]} \tag{10}$$

In each time step of the algorithm, each node selects a single neighbor at random and recieves a communication from that neighbor. The node updates its opinion according to the scheme outlined above in a seperate lise. Note in this algorithm each node is influenced by just one neighbor in each time step. In the previous implementations of the social impact theory, each node was influenced by *all* its neighbors in each time step.

3.2 Network structure

Network models are well known to be sensitive to initial conditions. In a purely random collection of network nodes, there would likely be little or no observable pattern of connection among the users. While the majority of existing models test their simulations on random graphs, this model will make use of a network that better approximates the structure of human social network. The data and information informing this network were derived from a study of a cell phone network aimed at uncovering the latent structure and organization of that network.

3.2.1 About the data and pre-processing

The data for this study was obtained from a major mobile and Voice Over Internet Protocol (VOIP) provider. The University of Virginia Institutional Review Board approved an annonmization process whereby the company encrypted user phone numbers by using a hashing function prior to providing the data set.

The data set contained 358,543,271 call detail records (CDRs) involving 50,855,844 unique users who placed or recieved calls on their network from June 1 through June 30, 2014. One day's records, those of June 3, were not used due to a data corruption issue. Other records were associated with non-standard users, such as service calls, telemarketing calls, voicemail calls, etc. Finally, certain recrods had indecipherable values, such as string entries for the area code. All such records were identified and omitted from this analysis.

This study was limited to call detail records in which at least one of the caller or callee had a number plan area code (NPA) within the Greater New York Metropolitan Area. This allowed for a more computationally tractable sample size. A potential trade-off is a loss of important information about the network structure. In the end, the analysis revealed interesting characteristics within the spatially-constrained network using a suite of network statistical measures which measured up well against similar studies performed on other data sets.

Following the aforementioned initial pre-processing of the CDR data were then constructed in a network formation in order to examine the node degree distribution. This graph contained a node for each unique user (caller and callee alike) with node attributes to keep track of the user's NPA, whether or not the user is identified as a Company user (based on the stream direction variable), and the user hash (as given in the original data set). The network was then populated with a directed edge for each unique call between caller and callee was with edge attributes capturing the stream direction associated with that call and a weight representing the number of times that unique call took place over the entire month.

The resulting graph was comprised of 4,465,313 nodes and 8,159,933 edges and a connected subgraph of 984,174 nodes and 3,312,927 edges. All further analysis will focus exclusively on the connected subgraph as unreachable nodes or clusters are uninteresting and have a negligible impact on the global opinion dynamics.

In this connected subgraph, the average weighted total (both in and out) node degree was 10.28 with a standard deviation of 214.72, a maximum of 238,349 and minimum of 1. The wide spread on this distribution was unrepresentative of average human social networks which is closer to a power law distribution. To approximate such a distribution nodes with degree less than 5 and greater than 500 were removed from the network. The resulting degree distribution is represented and approximated below.

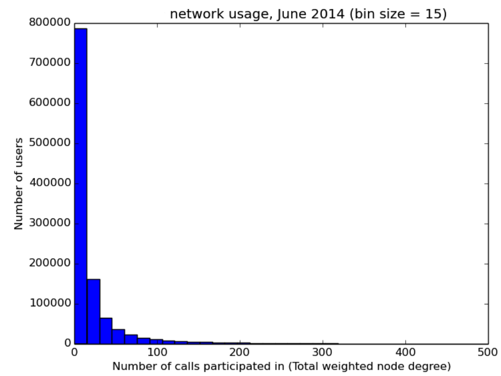

Figure 1: Weighted total node degree distribution.

This distribution can be approximated as a power law distribution. To verify this approximation, observe the distribution is nearly linear in log-log space.

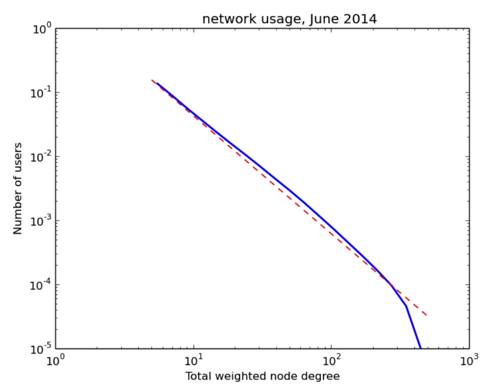

Figure 2: Weighted total node degree distribution in log-log space.

The power law distribution can be described as follows, where p is the number of users, x is the number of calls they participated in and α is the power law parameter.

$$f(x) \propto x^{-\alpha} = x^{-1.84270} \tag{11}$$

The power law nature of this distribution validates its suitability for approximating the structure of true social networks. All future analysis considers only the node degree constrained subgraph in which only nodes with degree between 5 and 500 are included.

3.2.2 Network statistics

The node degree constrained network contained 1,147,532 nodes and 2,935,739 edges with a connected subgraph composed of 522,186 nodes and 1,727,227 edges. To evaluate the latent structure of the network, a number of standard network statistics were computed and compared to statistical measurements of an Erdős-Rényi random graph. The random graph was generated by randomly swapping edges in the overall network (i.e., not the connected subgraph) according to an Erdős-Rényi distribution and then locating the connected subgraph within that randomized network. This approach is preferred to simply randomizing the connected subgraph because that subgraph is, in fact, an artifact of the larger network so randomizing the subgraph would preserve some of the characteristics of the data driven network which would not be truly random. The results of this analysis are summarized in the table below.

	Data	Random
Number of nodes	522,186	929,424
Number of edges	1,727,227	2,378,551
Average local clustering	0.00019	7.01341e-07
Global clustering	0.01664	5.69812e-06
Assortativity	0.01389	0.00033
Weighted average total degree	6.61537	5.11906
Average edge weight	4.60230	4.39491
Average shortest path length	Intractible	
Diameter	406.0	442.0
Is bipartite?	FALSE	FALSE

The number of nodes is the number of unique users (unique hashed user IDs) in each network. The number of edges is the number of unique calls in each network. Average local clustering is computed by iterating over each node and first determining its immediate neighbors (i.e., other nodes to which it is directly connected by one edge) and then counting the number of possible and actual connections among these neighbors. The local clustering coefficient for a node is given by:

$$C_i = \frac{\text{number of actual connections among neighbors}}{\text{number of possible connections among neighbors}} \tag{12}$$

The average local clustering coefficient is the average of the C_i taken over all nodes in the network graph. A graph is considered "small world" if its average local clustering coefficient is significantly higher than a random graph constructed with the same set of nodes and edges. In this case, the average local clustering coefficient for the data driven network is several orders of magnitude larger than the comparable random network. This indicates a "small world" clustering of nodes in the data driven graph.

The global clustering coefficient is computed by counting the number of closed and open triplets in the network. A triplet is a set of three nodes that are connected by either two or three edges; the former is "open" while the latter is called a "closed" triplet. The global culstering coefficient is computed as follows:

$$\bar{C} = \frac{\text{number of closed triplets}}{\text{number of open and closed triplets}} \tag{13}$$

The global clustering coefficient gives a measure of how densely connected a network is; the higher the global clustering coefficient, the more connected the network. In this case, the global clustering coefficient of the data driven network is several orders of magnitude larger than the comparable random network implying the data network is more densely connected.

Assortativity is a measure of homophily, the tendency for nodes, people in this case, networks to preferentially associate with nodes, or persons, of similar interests, tastes, etc. The assortativity coefficient is given by:

$$r = \frac{\sum_{xy} xy(e_{xy} - a_x b_y)}{\sigma_a \sigma_b} \tag{14}$$

Where $a_x = \sum_y e_{xy}$, $b_y = \sum_x e_{xy}$, and e_{xy} is the fraction of edges from a vertex of type x to a vertex of type y. σ_a and σ_b are the variances of a and b respectively.

In this case the characteristics x and y are node degree so assortativity measures the extent to which nodes connect by a rule of preferential attachment – the phenomenon in which popular nodes, simply by virtue of being popular, become more popular while less popular nodes become less popular. Observe that the assortativity coefficient is effectively the Pearson correlation coefficient of node degree between pairs of linked nodes. In this regard the data and random networks both demonstrate an assortativity coefficient close to zero which indicates some preferential process of node attachement.

The weighted average total degree in a network is the weighted average number of edges going in and out of a node. In this case, the weight is a factor representing the frequency of each unique call-pair. The average edge weight, then, is average number of times each unique call-pair occurs.

The average shortest path of a network is the average fewest number of steps required to go from any one node to another, called contact chaining.

$$\epsilon = \sum_{x,y \in V} \frac{d(x,y)}{n(n-1)} \tag{15}$$

Determining average shortest path is a computationally expensive operation and was intractable for such large networks, even when making use of the University of Virginia Advanced Computing Services.

The diameter of a network is the longest of all shortest paths in the network, given by:

$$\text{Diameter} = \max_{v \in V} \epsilon(v) \tag{16}$$

The significant difference in network statistics and measurements between the data driven graph and the random graph indicate that social networks have some inherent structure which cannot be replicated with a simple Erdős-Rényi random network. Girvan and Newman (2002) claim real-world networks typically have a clustering coefficient between 0.1 and 0.5. While the data driven network presented here has a clustering coefficient of 0.02, it is still much closer to the theoretical value given by Girvan and Newman (2002) than that of the random network. To put this in context, G. Palla and Vicsek (2005) present the following statistical properties of three benchmark real-world networks.

Network	Num. of nodes	Avg. node degree	Avg. clustering coeff.
Co-authorship	2,450	12.1	0.44
Word association	670	11.33	0.56
Protein interaction	82	1.54	0.17

The network studied here is a more appropriate simulation space for opinion dynamics models since telephone networks better represent the structure of information flow between people.

3.2.3 Network centrality

The final aspect of this network analysis involved the identification of nodes that were central to the flow of information in this network. Toward this end, two common centrality algorithms were applied. The first such algorithm, Page Rank, computes a score for each node v as follows:

$$PR(v) = \frac{1-d}{N} + d \sum_{u \in \Gamma^-(v)} \frac{PR(u) w_{u \to v}}{d^+(u)} \tag{17}$$

Where N is the total number of nodes, $\Gamma^-(v)$ are the in-neighbors of v, $d^+(u)$ is the out-degree of u, and d is a damping factor. In this case, the damping factor represents the likelihood of a node calling another node with whom it already has a connection (as opposed to randomly forging a new connection to a new node). The damping factor was set to the default $d = 0.85$.

The Page Rank algorithm considers both the quantity and quality of edges between nodes to estimate the importance of each node. The main assumption of the Page Rank algorithm is that nodes receiving more links from more important nodes are themselves important. All nodes were ranked according to this algorithm. A summary of the results is presented below.

The second centrality algorithm implemented, betweenness, considers – for each node – the number of shortest paths from all vertices to all others that pass through that node. A node with high betweenness centrality has a large influence on the network, under the assumption that connections follow the shortest paths. This algorithm is given by:

$$g(v) = \sum_{s \neq v \neq t} \frac{\sigma_{st}(v)}{\sigma_{st}} \tag{18}$$

Where σ_{st} is the total number of shortest paths from node s to node t and $\sigma_{st}(v)$ is the number of those paths that pass through v. The results of this algorithm are summarized below.

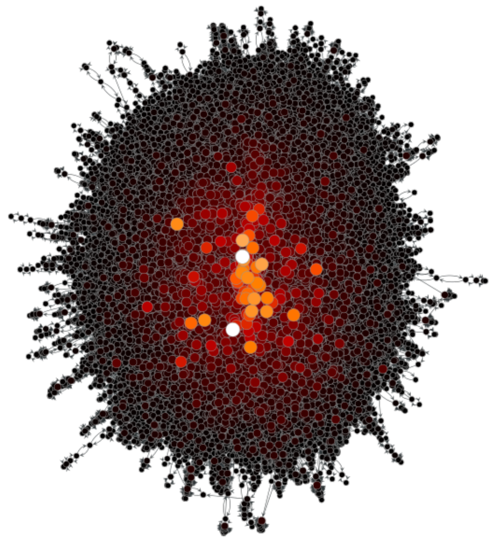

Figure 3: A rendering of the connected subgraph colored such that nodes with higher Page Rank are lighter in color.

Table 1: The 50 most central nodes according to the Page Rank algorithm.

User Hash	Page Rank	Network
0d17f0008374619a7a3f2d5b5d81be08	0.0001232926	Company
a2bd8e839f77a0bf4cc642249e1e7401	0.0001231742	PSTN
8bd3cbf9cb3916f6764494bc0aa11c2e	0.0001056989	PSTN
1b5f8e3261b18887558748342b519551	0.0001021668	Company
b5bb132f2cdf6e4382f77c1b2cf912d6	0.0001007798	Company
28c6c9862c922182431d247df733f9f6	0.0000978175	Company
b7970b2faa9905b3c7baafcef417f609	0.0000977702	PSTN
ff30ca3aaeaf8df16ffc7215e190318b	0.0000969494	PSTN
154d77a7b15305469a6d3c405f191022	0.0000954457	Company
15bbf72f80877700823856e650c9e117	0.0000949458	Company
28810315cad45c8a7cafb622381853bf	0.0000948474	PSTN
072d05f9346986146a6d3518c8fe0110	0.0000940846	Company
c74960d1c2fb81377124cd40b16f9411	0.0000938504	PSTN
9725aecf6a6aec216f3592160ff17695	0.0000921201	PSTN
fd0d2036495e3e32c566445ecd28eb76	0.0000916574	Company
02f16c6225bbc510c1ec211f98346075	0.0000914790	Company
39cc297d1309916849d7960211845659	0.0000909906	Company
199c0eca7bb9351b800f6703a0c9b166	0.0000898173	Company
aabf9f0d595303ff5b43418b23c511cb	0.0000894877	Company
a893d2913cf0148975189508de53865c	0.0000888918	Company
ba83bd2d66d2efc8feff90f09568e1e1	0.0000883434	PSTN
9d67ef6a95fa19cb16fc50b617a9c60b	0.0000881209	PSTN
e95bdd4a2f7a5d06128d902749e559e9	0.0000871566	PSTN
0bbed33b43a14e823755973a2e265d96	0.0000859067	Company
85350e77af17c39eae2efb334a5eecb5	0.0000831603	Company
2bba4780061b65afc6a6b8148497c510	0.0000814665	Company
c1e60f92bdbb380589b0e99e37635111	0.0000805223	PSTN
0bae1677620d08f95491a1dea467371f	0.0000799177	Company
cfbc738dfced5d81cf8bca189efa3c8a	0.0000795637	Company
9065bd065b79b0e3dac9535e21e4dfa2	0.0000792142	PSTN
38598f5fe3ad4ea25f0a6ac57efbb3fa	0.0000790604	Company
413db4519f48eb67b4344675698f898f	0.0000783035	PSTN
1d838185e7153965940ea1cfd137998c	0.0000778107	PSTN
0f5e194a1cd8f46c05c212a4b4394a6f	0.0000746046	PSTN
c9aaa09215500f13ce0f74e0f5c145dc	0.0000745676	PSTN
c9aaa09215500f13ce0f74e0f5c145dc:	0.0000745676	PSTN
359e456d6e8288f0e38fa8f034bad43a	0.0000744478	Company
7378e5252aa827a99f40415d422b411a	0.0000744001	Company
ba9343b963a3a7be344a9dc8e04b8636	0.0000722985	Company
fbbf7f2ce7b73c49a304b8fc30d690f0	0.0000710240	PSTN
e01d6524821d319391a84a1999e9d598	0.0000708263	Company
ff7ffe317dba633f9655e74275cbb2fe	0.0000708222	PSTN
c7b9005a8b4b4603bb54f4d35acc3b82	0.0000696782	Company
1068c6c1071d232b1a17f50933421e53	0.0000687694	PSTN
789709559bb9c8577065d7a539f3df07	0.0000686339	Company
3327ecbde5c47bbfc38121e19e8f62cd	0.0000685248	PSTN
a5a89586b85c65bc0329a0f8f2dad9e3	0.0000684130	Company
02e85843fead5ff2898ae8f8ba3462e4	0.0000681548	PSTN
aaf5b14fc65cf75dc9d7d552a94f317e	0.0000673218	Company
7a5abe0a2b93b075f63835de74b035a8	0.0000672901	Company

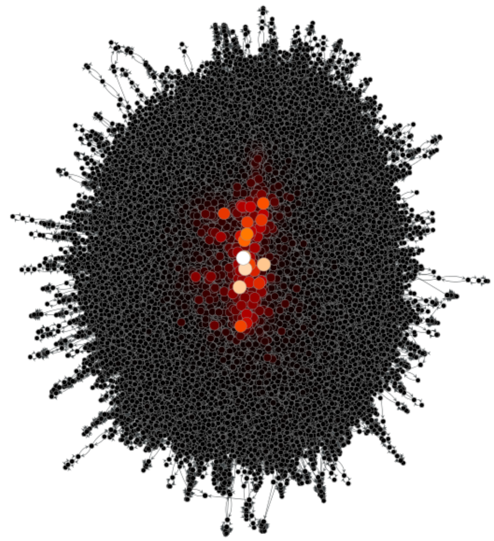

Figure 4: A rendering of the connected subgraph colored such that nodes with higher betweenness measure are lighter in color.

Table 2: The 50 most central nodes according to the betweenness algorithm.

User Hash	Betweeness	Network
0344c47c1a68f833be71636e1cbde4b3	0.0135690187617	Company
a893d2913cf0148975189508de53865c	0.0124742430163	Company
cbd41c6103064d3f0af848208c20ece2	0.0122196391687	Company
b5bb132f2cdf6e4382f77c1b2cf912d6	0.0120588984679	Company
a2bd8e839f77a0bf4cc642249e1e7401	0.0112880322728	PSTN
a2a6bb076324e18c34d0ba4dc64374f6	0.0099676036100	PSTN
1b5f8e3261b18887558748342b519551	0.0094149414751	Company
7e76b1118b52dbcd0a7e4d0157ab1c25	0.0089498381588	PSTN
6db76106f006028acf2fb7188328161e	0.0085825884342	PSTN
ff7ffe317dba633f9655e74275cbb2fe	0.0085018232728	PSTN
aad4e440185da8f2bf109e7fe7aa6925	0.0082357435921	Company
1e8dd4e060ad940cc734b50d1a9c59d4	0.0082256606363	Company
a849add9a28b95ae1050dd131f99ae04	0.0081315344560	PSTN
9725aecf6a6aec216f3592160ff17695	0.0078536251766	PSTN
5794f081a24d24174653895847ff7829	0.0076461099262	Company
c44e781a0f7dc2f6a5a6987a2bbac34a	0.0071998484203	Company
81bf2141b4c8278e54e17861d82c305d	0.0071638976580	Company
db0096ebc89a7ecf04ddb34ec51f5f02	0.0070589001844	PSTN
ccaff5445e4fd2c46d12706ceb9b7dc3	0.0067832770594	Company
484edae9268e91b7e58ce78a46e4712a	0.0066931397039	PSTN
342f652451c195605261cf7af3aeac0b	0.0065678370666	PSTN
815de431bf7d4b93f3fbe81716002eb2	0.0065211696319	Company
8bd3cbf9cb3916f6764494bc0aa11c2e	0.0063671935260	PSTN
2e0c16b782b0fb86fd172bfea67ccffb	0.0062983084577	Company
d846e0c47e2598424e458b985fe04874	0.0062978285449	Company
08196c9c5e8959d290c8d4deb3bc24cf	0.0062329592982	Company
81cc6aae316673520cf0e9750c559fe2	0.0061986193370	Company
c250b25958c7f19496d9b685d6eb9c84	0.0058514522066	Company
9d67ef6a95fa19cb16fc50b617a9c60b	0.0056424915174	PSTN
e160a366bd740281c409875ab8f0dde4	0.0056046530625	Company
948cf108a182c90d7ac1dff66d7c9500	0.0055921480811	Company
292c4e9124aa0eaa56081fcf4bc6a9b9	0.0054749444740	Company
d49a2c7f43fdb95853c72abebafad9b7	0.0054524892120	Company
ce93871977318c8b005a3d683aec7af4	0.0054507515795	Company
5e11d09069e341a794b831b5bbc43dc2	0.0054501607979	Company
9b961213d45a21d751b970bfb2659217	0.0053884252201	Company
fd9a673b0d2e04e6a680d84c86ab5360	0.0053564240533	Company
ce9231d1b5a455179e32bdceb513eee4	0.0053094520566	PSTN
62ed08a6afda4066b89d03ef092d95b7	0.0051158401254	Company
bcf929a7063e5ee52cf8df05d1978501	0.0049790668304	PSTN
108c54a6d36c70caf86f331fdf60c77d	0.0048986667175	Company
74114357629e7abf7e358fc019a097e2	0.0048423302350	Company
db408cbc8142c83a0669a60ac379e696	0.0048414023343	Company
db408cbc8142c83a0669a60ac379e696	0.0048414023343	Company
2c92ad05ff2528a4dd10b840022dbc3e	0.0048243487380	PSTN
64f820a1f16258854dcbcfaabd28198f	0.0048240482222	Company
d337d88fc4fee40c587e6b68a3e79030	0.0048065858148	Company
5d9334f5b0457bed11217692194335fe	0.0047894309361	PSTN
d610877735094573a21f3e2050398a76	0.0047600080273	Company
81361ddd22908aad6bbf5135260bda57	0.0047487310325	Company

3.3 Simulation methods

The social impact thoery and variations described above were implemented on a social network with similar structure to the graph formulated above. The data driven graph would have been simply too large for use under simulation and may have caused several computations to be unnecessarily expensive. For a more convenient study, a new graph was constructed with the same average local clustering coefficient, global clustering coefficient, and assortativity. Because the powerlaw nature of the node degree distribution, demonstrated above, implies the data driven network has some scale-free characteristics, taking a smaller scale replica of the data driven graph still preserves its structural integrity.

This model aims to simulate the dynamics of word-of-mouth influence. Word-of-mouth relationships are not permanent. To simulate this, at the end of each iteration a double edge swap is performed over a specified number of edges. A double edge swap removes two randomly chosen edges $u \longleftrightarrow v$ and $x \longleftrightarrow y$ and creates the new edges $u \longleftrightarrow x$ and $v \longleftrightarrow y$. If either the edge $u \longleftrightarrow x$ or $v \longleftrightarrow y$ already exist no swap is performed so the actual count of swapped edges may be less than the specified number of swaps. After this operation, new values for distance d_{ux} and d_{vy} are generated.

In the follow section, we will make use of Monte Carlo simulation methods to explore:

1. The global and local dynamics of the social impact theory as applied to the network model and in the manner detailed above.

2. The impact of key aspects and parameters on the model outcomes:

   – Varying α.

   – Attributing values other than zero to the random field h_i.

   – Investigating nodes with highest betweenness as predictors for global or local behavior.

3. The global and local dynamics of the three variations on the social impact theory model detailed above (memory, three opinion state space, and Bayesian updating). In the case of the memory model, three further variations will be presented. When stating a discrete opinion the kind of point estimate may be:

   – The mean of the memory distribution.

   – The mode of the memory distribution.

   – A random draw from the memory distribution with varying emphasis on more salient, or more recent, memories.

4 Results

4.1 Benchmark results

The benchmark social impact theory detailed in section 3.1.1 was tested on various randomly generated networks in accordance with the simulation methods outlined above. Two main classes of global dynamics, steady and unsteady, were observed with several specific cases in the first class. In the first stead case, the initial opinion distribution was about equal between $\sigma = +1$ and $\sigma = -1$. An example of such a simulation is presented in Figure 5.

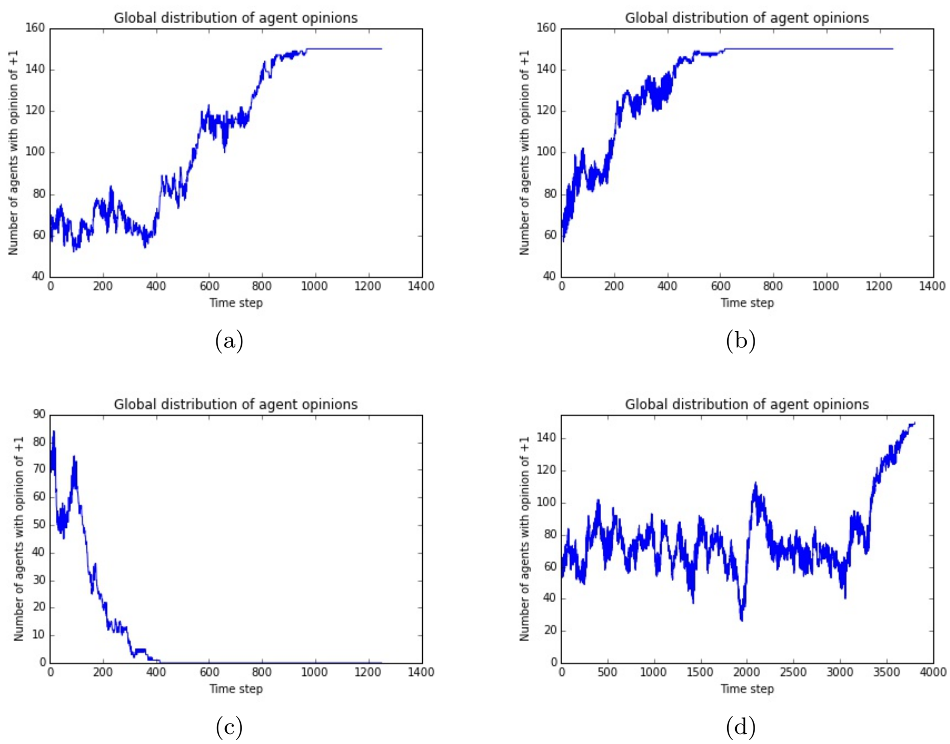

Figure 5: Representative results of the benchmark social impact theory model.

Figure 5 (a) presentes the global opinion dynamics over time with an initial distribution around 50%. Of 150 agents, inititally 74 agents held an opinion of $\sigma = +1$ while 76 held $\sigma = -1$. In this example, the global dynamics show some variability within few hundred time steps but the general trend is clearly approaching consensus.

In the second case of steady global dynamics, the initial opinion distribution favored either $\sigma = +1$ or $\sigma = -1$. Examples of such simulations are presented in Figures 5 (b) and 5 (c). Figure 5 (b) presents the global dynamics over time of a simulation with an initial distribution favoring agents with opinion -1. Of 150 agents, initially 63 agents held an opinion of $\sigma = +1$ while 82 held $\sigma = -1$. In this case, consensus is acheived in approximately 600 iterations. In Figure 5 (c), the initial distribution favored agents with opinion $+1$. Of 150 agents, initially 87 agents held an opinion of $\sigma = +1$ while 68 held $\sigma = -1$. In this example consensus is achieved within approximately 400 iterations. Note that in both cases the initial minority opinion grew to dominate the entire network. Likewise, in both examples convergence was achieved within a few short time steps.

The second main class of benchmark results is defined by relatively unstable global dynamics. In comparison to Figures 5 (a), 5 (b), and 5 (c), observe the noticeable difference in the globaly dynamics over the first 1200 iterations in Figure 5 (d). This class of models requires significantly more time steps to reach convergence. In this example, consensus is acheived at approximately 4000 iterations. This particular simulation offers an opportunity to examine the impact of different parameters on the model.

4.1.1 Varying α

By definition, α adjusts the relative impact of distance, d_{ij}, on social impact. In order to examine the impact of the parameter α on the global dynamics, the initial network presented in Figure 5 (d) above is reproduced and controlled in ever aspect, including edge mutations, while varying α. In all the benchmark model shown, $\alpha = 5.5$.

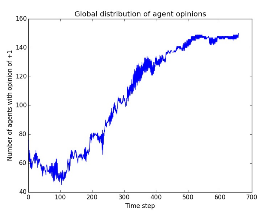

(a) $\alpha = 2.5$

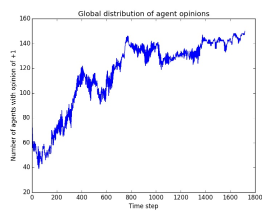

(b) $\alpha = 4$

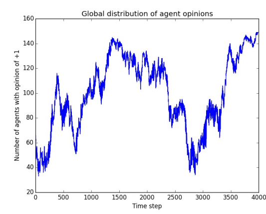

(c) $\alpha = 5$

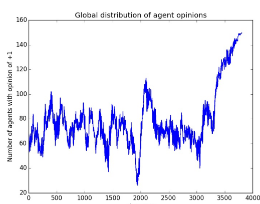

(d) $\alpha = 5.5$

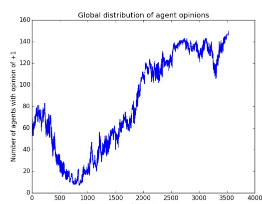

(e) $\alpha = 7$

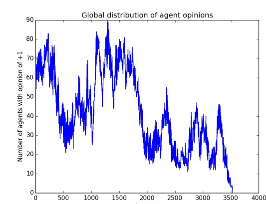

(f) $\alpha = 8$

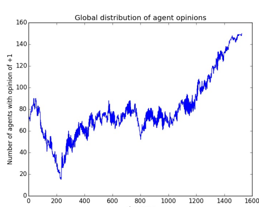

(g) $\alpha = 10$

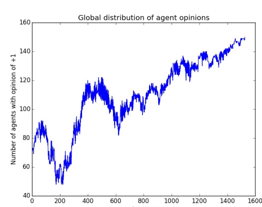

(h) $\alpha = 15$

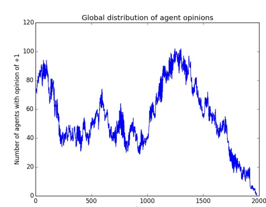

(i) $\alpha = 20$

Figure 6: Varying α for the system presented in Figure 5 (d).

This model is very sensitive to α in terms of time required to reach consensus and which opinion dominates. For $\alpha = 2.5$ and $\alpha = 3$, the network reaches consensus in about 700 time steps but in the former case converges on $\sigma = +1$ and in the latter on $\sigma = -1$. When $\alpha = 3.5$ or $\alpha = 4$, the systems is relatively stable and converges to $\sigma = +1$ in about 1400 time steps. For $4.5 < \alpha < 12$, the system becomes considerably less stable and identifying a pattern in among the different values for α is difficult.

When $\alpha = 12$ and $\alpha = 13$, the global dynamics are very similar, converging on $\sigma = -1$ in about 650 time steps. Yet, when $\alpha = 14$, the global dynamics are significantly different – requiring over 200 time steps to convergence. The system behaves nearly identically when $\alpha = 15, 16$ – converging in about 1600 time steps to $\sigma = +1$. Likewise, when $\alpha = 17, 18, 19, 20$ the global dynamics are nearly identicaly, converging in about 2000 time steps to $\sigma = -1$.

4.1.2 Varying h_i

Now, we introduce nonzero values for the random field h_i which represents all forms of influence on opinion other than dynamic social influence. In each trial, nodes were randomly assigned a value h_i from the specified uniform distribution. This random field may represent influence from a personal bias, a static social influence, an economic influence, among others. Again, these simulations were implemented on the same graph presented in Figure 5 (d) above.

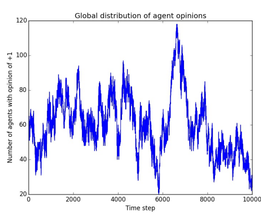

(a) $h_i = (-0.5, 0.5)$

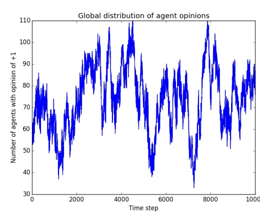

(b) $h_i = (-1, 1)$

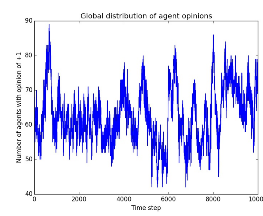

(c) $h_i = (-5, 5)$

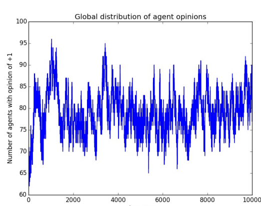

(d) $h_i = (-10, 10)$

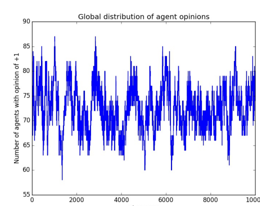

(e) $h_i = (-25, 25)$

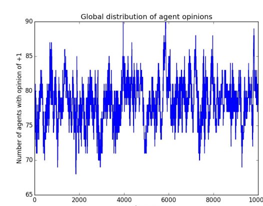

(f) $h_i = (-50, 50)$

Figure 7: Global dynamics for various values of h_i for the system presented in Figure 5 (d).

For computational purposes, the number of iterations was limited to 10,000. For $h_i = (-0.5, 0.5)$, the system is nearly converging on $\sigma = -1$ while for $h_i = (-1, 1)$ the system is nearly converging on $\sigma = +1$. While it is unclear towards which opinion it will converge, the system is converging for $h_i = (-5, 5)$. On the other hand, for $h_i = (-10, 10)$, $h_i = (-25, 25)$, $h_i = (-50, 50)$, and $h_i = (-100, 100)$ the system becomes chaotic with no discernable pattern to the global dynamics.

4.1.3 Betweenness centrality

Now, we apply the betweenness centrality algorithm as detailed in the section 3.2.3. Using the results of a simulation presented in the following figure as a benchmark, we examine whether the global behavior can be predicted by tracking the behavior of the nodes in the top 10% of betweenness centrality.

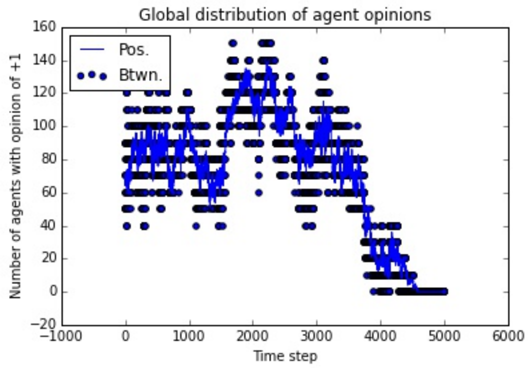

Figure 8: Benchmark simulation to examine the predictive potential of betweenness centrality. The scatter plot points for betweenness represent the number of agents with the highest 10% betweenness centrality whose opinion is +1. This count is then scaled by 10 in order to visualize on the same scale as the overall network counts.

The simulation presented in Figure 8 requires about 4500 iterations to reach consensus. At such a scale it is difficult to asses the predictive potential of the betweenness centrality algorithm. So we will examine the graph at various intervals, each of ten time steps.

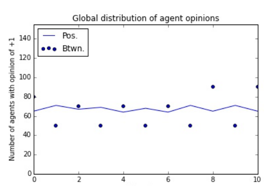

(a) Time steps 0–10.

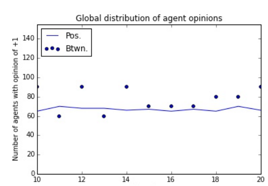

(b) Time steps 10–20.

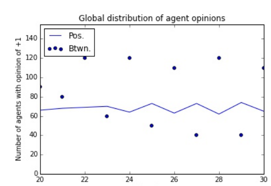

(c) Time steps 20–30.

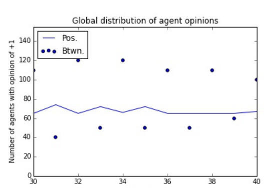

(d) Time steps 30–40.

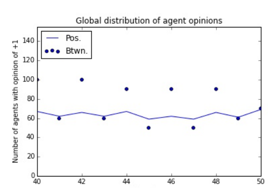

(e) Time steps 40–50.

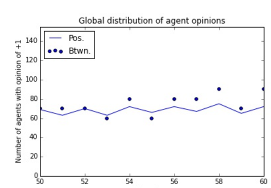

(f) Time steps 50–60.

Figure 9: Local behavior of nodes with highest betweenness against global behavior.

At the beginning of the simulation when the systems undergoes considerable variability, the local behavior of the nodes with highest betweenness consistently preempts the global behavior of all nodes by about one time step. This is easily noticeable in Figure 9 (c) when, in each time step, the nodes of highest betweenness take a certain position and in the very next time step, the global behavior follows suit.

This framework may be used to ask more interesting questions. In one time span during the simulation, the global behavior nears consensus at the positive opinion but does not converge to that position. A closer look at the global and local dynamics shows that the nodes with highest betweenness were consistently drawing the global system away from that convergence. This correlation is not surprising and, in fact, confirms the effectiveness of the betweenness algorithm at selecting the most central or influential nodes in a social network.

4.2 Social impact theory with memory

Now, we implement agent memory in the manner detailed in section 3.1.2. Figure 16 presents a new benchmark model from which all other aspects of the simulation will be controlled.

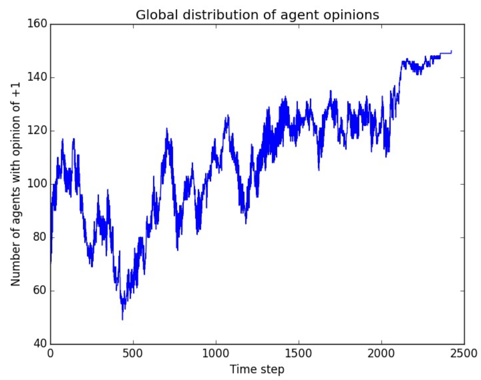

Figure 10: Benchmark simulation (i.e., without memory) used to control the experiment for a social impact model with memory.

The results of the updating rules with memory are presented in Figure 11 below. Observe that a random choice from an agent's recent memory, however defined, produces a global dynamic not too dissimilar from the benchmark. However, a random choice from an agent's entire memory leads to a chaotic global dynamic with no reasonable indication towards convergence or patterned behavior. This is not surprising as a random draw from the entire distribution will have no reinforcement towards or against a particular opinion and thus no chance of consensus.

It is no surprise that, in a two opinion state space with $\alpha = \pm 1$ the mean (rounded mean) and mode of a distribution are almost always symmetrical. As a result, Figures 11 (e) and 11 (f) display the same global dynamic. In this case, the system converges in less than 600 iterations, as opposed to the nearly 2500 required for the benchmark system to converge.

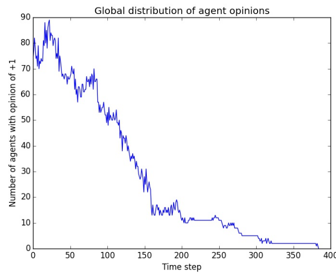

(a) Random choice from 5 most recent interactions.

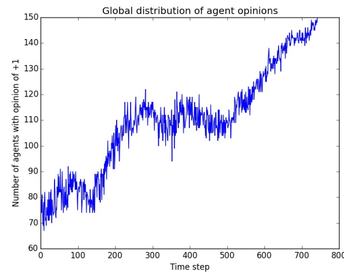

(b) Random choice from 15 most recent interactions.

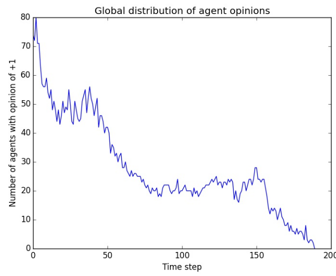

(c) Random choice from 100 most recent interactions.

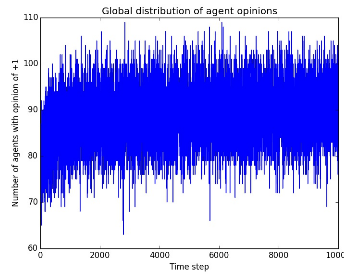

(d) Random choice from all interactions.

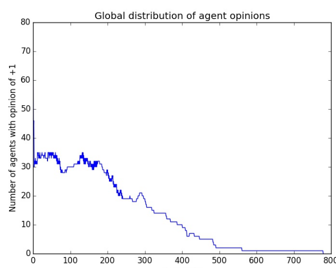

(e) Choose the mode of the memory distribution.

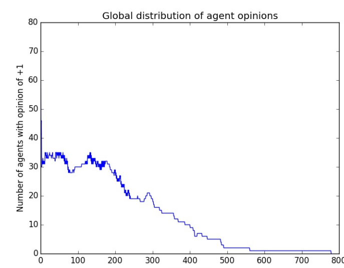

(f) Choose the mean of the memory distribution.

Figure 11: Simulation results for various interpretations of agent memory.

4.3 Social impact theory with three opinion state space

Considering a third opinion state, neutrality, leads to significantly different global dynamics. Figure 12 presents four representative results of such simulations. In each case, a clustering of opinions is randomly assigned to the network and such clustering largely persists, albeit with some fluctuations and patterned variation. In simulations with a two opinion state space, it is nearly impossible for more than one opinion to persist. However, with the introduction of some measure of apathy or ambivalence, agents' social influence is diluted and are able to better preserve their initial opinions.

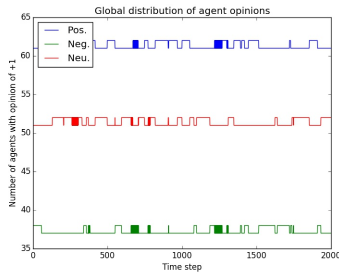

(a)

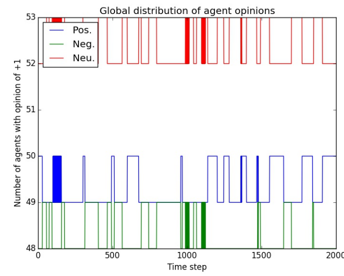

(b)

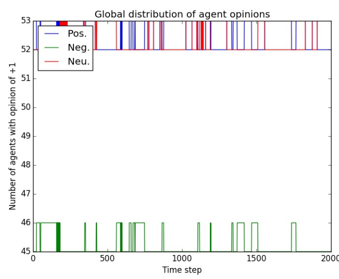

(c)

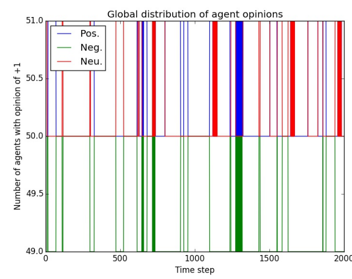

(d)

Figure 12: Representative global dynamics for simulations with three opinion state space.

4.4 Social impact theory with Bayesian updating

The three opinion state space, as presented in the previous section, allowed for clustering and the coexistence of multiple opinions *ad infinitum*. The Bayesian updating procedure allows for the coexistence of multiple opinions with greater variability in the global dynamics. Because agents are attempting to identify the opinion with the highest likelihood, they can be more readily influenced and yet still maintain their independence. This leads to very interesting global dynamics as demonstrated in Figure 13 below. In some cases, one opinion will still dominate the longterm dynamics but in most cases there will be some coexistence of the three opinions in the longterm.

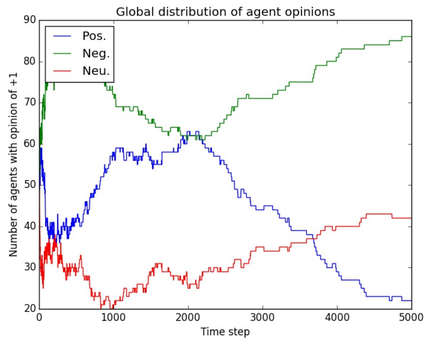

(a)

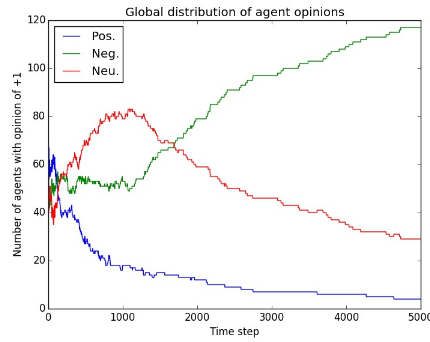

(b)

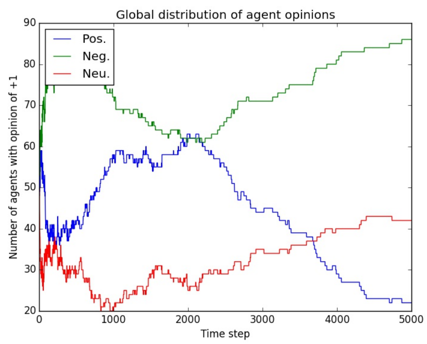

(c)

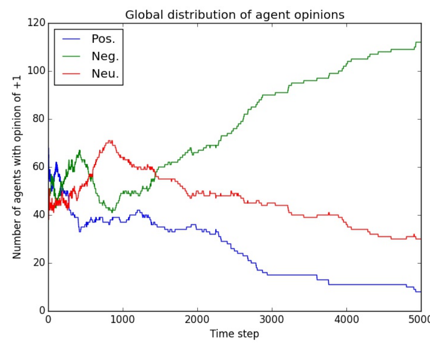

(d)

Figure 13: Representative global dynamics for simulations with Bayesian updating.

5 Discussion

This project aimed primarily to develop more interesting and realistic agent rules for opinion dynamics models and secondarily to apply these rules to a more representative social network model. Based on the social impact theory, this project arrives at a Bayesian method of updating that allows for a rich coexistence of more than two opinion states over long periods of time, often limitlessly. This is unique among opinion dynamics models and has not been demonstrated in the realm of social impact theoretic models.

The vast majortiy of opinion dyanmics models presented elsewhere that demonstrate some clustering or coexistence of multiple opinions require the imposition of stubborn agents in the network around which such clusters will form. Presented here is a model in which clustering occurs organically as a result of the updating procedure. This is far more representative of the dynamics observed in the natural world.

In arriving at this model, the effect of various parameters on the benchmark social impact theory were examined. It was shown that adjusting α, a tuning parameter of the rate at which influence changes with disctance, alters the global dynamics in a way that is not fundamentally different than arbitrarily chose α. Furthermore, varying h_i, a measure of all influence on agent opinion other than social impact, can increase variability in the system. But, as with α, h_i, does not highlight any esential characteristics of agent opinions and behavior.

A model of social impact with agent memory was implemented with various methods of operating over memory at the agent level. Surprisingly, this was not a very rich adaptation on the model. In the end, agent memory simply caused the global dynamics to be either more or less chaotic depending the implementation but did not offer much in the way of a more fruitful model.

At first glance, the social impact theory model adapted to operate over a three opinion state space seems interesting. A clustering and coexistence of opinions is observed in the global dynamics. But a closer look at the local behavior shows that less than 5% of nodes are actively changing their positons over time while the majority of nodes retain their initial opinion.

It is, finally, the social impact theory model with Bayesian updating that offers a rich opinion dynamics model in which agent behavior is turly nonlinear. However, there does remain some additional work to examine more thoroughly the local dynamics of this model. It was shown here that the betweenness centrality algorithm may identify nodes whose local behavior preempts global behavior. But that is far from a complete description of the local dynamics and may not offer any *predictive* power, that is, while this project was meant, at the outset, to consider questions of causality, in the end such questions must be left for future study.

Still, this model offers a fertile experimentation ground for those interested in political opinions, product adoption, or information diffusal. A more sophisticated model may consider grounding the parameters, p_i, s_i, and a_i, in some empirical or at least deliberate methodology. For example, the American National Election Survey routinely measures parameters such as political awareness, political apathy, and party affiliation that may be considered to give this model further grounding in much the same way it was tested here on a data-drive network topology.

6 Acknowledgements

I would like to thank Professor Gerard Learmonth, my advisor, for his endless support and guidance during the undertaking of this project. I would also like to thank my thesis committee chair and member, Professor Alfredo Garcia and Professor Paul Humphreys, for their invaluable feedback. Finally, I gratefully acknowledge support for this project from the Data Sciences Institute and the Department of Systems and Information Engineering.

7 Bibliography

R. Axelrod, "The dissemination of culture: A model with local convergence and global polarization," *Journal of Conflict Resolution*, vol. 41, no. 2, pp. 203–226, 1997.

C. Castellano, S. Fortunato, and V. Loreto, "Statistical physics of social dynamics," *Reviews of Modern Physics*, vol. 81, no. 2, pp. 591–646, 2009.

M. H. DeGroot, "Reaching a consensu," *Journal of the American Statistical Association,*, 1974.

B. Fotouhi and M. G. Rabbat, "The effect of exogenous inputs and defiant agents on opinion dynamics with local and global interactions," *IEEE Journal of Selected Topics in Signal Processing*, 2013.

N. E. Friedkin and E. C. Johnsen, "Social influence and opinions," *Journal of Mathematical Sociology*, 1990.

I. F. G. Palla, I. Derenyi and T. Vicsek, "Uncovering the overlapping community structure of complex networks in nature and society," *Nature*, vol. 435, 2005.

M. Girvan and M. E. J. Newman, "Community structure in social and biological networks," *Proceedings of the National Academy of Sciences*, vol. 99, no. 12, pp. 7821–7826, 2002.

R. Hegselmann and U. Krause, "Opinion dynamics and bounded confidence models, analysis, and simulation," *Journal of Artifical Societies and Social Simulation*, 2002.

R. A. Holley and T. M. Liggett, "Ergodic theorems for weakly interacting infinite systems and the voter model," *The Annals of Probability*, 1975.

D. Kempe, J. Kleinberg, S. Oren, and A. Slivkins, "Selection and influence in cultural dynamics," *ACM EC*, 2013.

A. Koulouris, loannis Katerelos, and T. Tsekeris, "Multi-equilibria regulation agent-based model of opinion dynamics in social networks," *Interdisciplinary Description of Complex Systems*, 2012.

B. Latane, "The psychology of social impact," *American Psychologist*, vol. 36, no. 4, pp. 343–356, 1981.

Q. Li, L. A. Braunstein, H. Wang, J. Shao, H. E. Stanley, and S. Havlin, "Non-consensus opinion models on complex networks," *Journal of Statistical Physics*, 2012.

Y. Liu, X.-M. Si, and Y.-C. Zhang, "Comparing effects of cluster-coupled patterns on opinion dynamics," *International Journal of Modern Physics C*, 2012.

A. C. R. Martins, "Continuous opinions and discrete actions in opinion dynamics problems," *International Journal of Modern Physics*, 2007.

——, “Mobility and social network effects on extremist opinions,” *Physics Review E*, 2008.

——, “Bayesian updating as basis for opinion dynamics models,” *XI Brazilian Meeting on Bayesian Statistics*, 2012.

J. Shao, S. Havlin, and H. E. Stanley, “Dynamic opinion model and invasion percolation,” *Physical Review Letters*, 2009.

K. Sznajd-Weron, “Sznajd model and its applications,” *ACTA PHYSICA POLONICA B*, vol. 36, no. 8, pp. 2537–2547, 2005.

F. Xiong, Y. Liu, and Z. Zhang, “Dynamics with co-evolution of individual inclination and opinion,” *International Journal of Modern Physics C*, 2010.

H.-X. Yang, W.-X. Wang, Y.-C. Lai, and B.-H. Wang, “Convergence to global consensus in opinion dynamics under a nonlinear voter model,” *Physics Letters A*, 2012.

E. Yildiz, D. Acemoglu, A. Ozdaglar, A. Saberi, and A. Scaglione, “Discrete opinion dynamics with stubborn agents,” *Operations Research*, 2011.

S. Zhang, L. Chen, D. Hu, and X. Sun, “Impact of opinions and relationships coevolving on self-organization of opinion clusters,” *Interdisciplinary Description of Complex Systems*, 2013.