Opinion Formation on Networks

Alexander Pan*

Supervisors: Heather Z. Brooks†, Franca Hoffmann‡, Mason A. Porter§

Abstract

People are increasingly reliant on online media sources for information, and media outlets continue have a great impact on public opinion. With the rise of extremism in society, it is especially important to understand how media outlets cause polarization of ideologies. Mathematical models can give us new insight into the dynamics of opinion formation and propagation. One popular type of opinion model is a bounded-confidence model. In such a model, agents start with some opinion and evolve their opinion by interacting with other agents whose opinion is sufficiently similar to their own. A recent extension of the model captures the impact of media on opinions through a network; the network mathematically characterizes the population by representing both individuals and media outlets and the information channels between them, while imposing constraints on who interacts with whom. In certain models, individual opinions will eventually converge to a steady state. Currently, little is known about the distribution of opinions at steady state and how the media influences this distribution. We examine bounded-confidence models on synthetic networks with various levels of media influence. In a complete network, we observe numerically that there is a phase transition, indicating a sudden change in qualitative behavior, from a single opinion group to several opinion groups. In ongoing research, we are working on a characterization and rigorous proof of the observed bifurcation and analyzing variations of the model.

1 Introduction

Media outlets and social media enjoy a widespread following in U.S. society [17], and they also affect public opinion [1]. Because of the growing political polarization in society, we need to gain a better understanding of how media influences opinion [10]. To make this question mathematically tractable, there are three issues that must be addressed. First, we must determine how to model social relationships in a population. A natural representation of a population and its social relationships is a network, where nodes represent individuals carrying opinions that are measured on a continuous scale, and edges represent relationships between nodes, allowing social interaction. Second, we have to model how individuals in the population interact (influence each others' opinions). In this report, we use an opinion model to represent social interactions. Based on our current understanding, social dynamics are not governed by any clear laws, so our choice of opinion model is based on both mathematical tractability and observation. One possible opinion model is the DeGroot model [9], where individuals update their opinion based on some linear combination of all other opinions. The model we study in this report is a variation of a DeGroot model on a network. One such model is a bounded-confidence opinion model, which was first suggested by Deffuant *et al.* in [8] and Hegselmann and Krause in [15]. Third, we need to come up with some way of measuring media impact on society. The goal of this SURF project is to study measures of media impact on bounded-confidence opinion models, extending the work done by Brooks and Porter in [2].

1.1 Modeling framework

We now give an introduction to the model we started with, taken from [2]. From now until noted, we will mostly follow the presentation of [2], a recent work that introduced media nodes into bounded-confidence

*Department of Computing and Mathematics Sciences, Caltech.

†CAM Assistant Professor, Department of Mathematics, UCLA.

‡Von Karman Postdoctoral Instructor, Department of Computing and Mathematical Sciences, Caltech.

§Professor, Department of Mathematics, UCLA.

models and examined the effects of doing so. Consider a social network represented by a directed graph $G = (V, E)$, where the set V of nodes encodes accounts and the set E of edges encodes followership, such as on Twitter. A directed edge from j to i is in the set E if account i *follows* account j. This allows us to build an adjacency matrix $\mathbf{A}$, with $\mathbf{A}_{ij} = 1$ if i follows j and $\mathbf{A}_{ij} = 0$ otherwise. In general, followership is not reciprocal, so $\mathbf{A}$ may not be symmetric. We also assume that $\mathbf{A}_{ii} = 1$ for all i, so each account follows itself.

We suppose that there are N individual accounts in a network, and we suppose that the network also has M media accounts. These media accounts produce content and do not change their ideology, so each media account has only out-edges. For our purposes, we assume that each media account has n_M followers. After we include media nodes, there are a total of $N + M$ nodes in our network.

To capture the behavior of the network over time, we give each node i (where $i \in \{1, \ldots, N + M\}$) an opinion $\mathbf{x}_i^k$ at discrete time k, where we take $\mathbf{x} \in [-1, 1]^d$. For simplicity, we take $d = 1$. In this report, we interchange the words "ideology" and "opinion" freely. For $d = 1$, we can imagine the opinion space to be political ideologies (we can still use the analogy of political ideologies if $d > 1$ by allowing more nuanced ideologies), with $\mathbf{x} = 1$ representing a strongly conservative opinion, $\mathbf{x} = -1$ representing a strongly liberal opinion, and $\mathbf{x} = 0$ representing a moderate opinion.

Each node in our network updates its ideology synchronously at each discrete time step k by averaging the opinion of all accounts to which it is receptive. Although accounts may follow many other accounts, we say that an account i is *receptive* to account j if i follows j and $\mathbf{dist}(\mathbf{x}_i - \mathbf{x}_j) < c$, where $\mathbf{dist}$ is some notion of distance between i and j in the opinion space. We use c as a notion of how similar the opinions of two accounts have to be in order for the accounts to influence each other. A smaller c means fewer accounts will be able to influence any given account. Note that this definition of receptiveness implies that any media node is not receptive to any node except itself (because each media node follows only itself). Intuitively, account i is only receptive to (and hence can be influenced by) account j if their ideologies are sufficiently "close" (determined by c) and i receives j's content. This captures the assumption that individuals are influenced only by accounts who have a relatively similar opinion [13]. The idea of an opinion cutoff comes from bounded-confidence opinion models [19]. In this class of models, agents carry opinions and evolve them by interacting with other agents whose opinion is sufficiently similar to their own. Bounded-confidence models use confidence intervals; for an opinion $\mathbf{x}$, a confidence interval is an interval of opinions centered on $\mathbf{x}$ that can influence $\mathbf{x}$.

For $d = 1$, we use the ℓ_2 norm to measure distance. That is, we define $\mathbf{dist}(x, y) = |x - y|$, where $|\cdot|$ denotes Euclidean distance. We can then define the set of all accounts that are receptive to account i as $I_i = \{j \in \{1, \ldots, N + M\} \mid \text{account } i \text{ is receptive to account } j\}$; equivalently, we can write $I_i = \{j \in \{1, \ldots, N + M\} \mid \mathbf{A}_{ij} = 1 \text{ and } \mathbf{dist}(\mathbf{x}_i, \mathbf{x}_j) < c\}$.

We define the update rule for each account as follows:

$$\mathbf{x}_i^{k+1} = \frac{1}{|I_i|} \sum_{j \in I_i} \mathbf{x}_j^k = \frac{1}{|I_i|} \sum_{j=1}^{N+M} \mathbf{A}_{ij} \mathbf{x}_j^k \eta \left(\left| \mathbf{x}_i^k - \mathbf{x}_j^k \right| \right) , \quad \text{for } k \in \mathbb{N} , \quad i \in \{1, \ldots, N + M\} , \tag{1}$$

where

$$\eta(x) = \begin{cases} 1 , & \text{if } x < c , \\ 0 , & \text{if } x \geq c . \end{cases}$$

Note that each media account i, with $i \in \{N + 1, \ldots, M\}$, has $I_i = \{i\}$, so the media account ideologies are constant over time.

To examine the long-term behavior of our model, we simulate the model numerically on various types of networks. For each trial, we create a network and fix certain aspects, such as the number of nodes or the type of network. We first construct a network for the non-media nodes and then add media nodes to the network. For the network for the non-media nodes, we use either a Watts–Strogatz network defined in Section 1.2 (since a Watts–Strogatz network is undirected, to turn it into a directed network, all edges are made reciprocal), a random network model with short mean path length, or a complete network, a network where all nodes are adjacent. We assign each node in the network an opinion that we draw from a uniform distribution over our opinion space. In all of the numerical simulations, we observed that if we allow the opinions on the network to evolve over time according to (1), we eventually reach a point in time at which the opinions do not change within a tolerance interval of width ϵ. We use the word *convergence* to describe this situation. Mathematically, we define "convergence" for a trial to occur when $\mathbf{dist}(\mathbf{x}_i^{k+1}, \mathbf{x}_i^k) < \epsilon$ for all i for ten consecutive time steps.

1.2 Network definitions

In this report, we focused on studying complete networks and Watts–Strogatz networks; we briefly looked into star networks and ring-lattice networks [20].

- A *star network* is a deterministic network with a special node (known as the "star") that shares an edge with every other node; any non-star node has only one edge, which it shares with the star.
- A *ring lattice* is a deterministic network with two parameters: the number of nodes N and the degree k. For each node i in $1, \dots, N$, the node i shares an edge with nodes $i+1, \dots, i+k$ (where node $N+1$ is equivalent to node 1), so each node has degree k.
- A *complete network* is a deterministic network, where all pairs of nodes share an edge.
- A *Watts–Strogatz network* is a random graph generated from three parameters: the number of nodes N, the mean degree k, and the rewiring probability β. We start with a ring lattice with N nodes and mean degree k. Then, for each i in $1, \dots, N$ and for each j in $i+1, \dots, k$, the edge (i, j) is rewired with probability β. Rewiring means that (i, j) is replaced (i, k), while avoiding self-loops and duplicating edges.

The above networks are all undirected. To make them directed, for each edge (i, j) in the undirected version of the network, we add the edge $i \to j$ and $j \to i$ in its directed counterpart.

1.3 Measuring media impact

We look at a network at convergence to examine the effects of media on the opinions of the individuals over time. Each media node has n_M followers, which are selected uniformly at random from the non-media nodes. Although (1) works for all dimensions d, we use $d = 1$ in our work. For simplicity, we set all of the media nodes to have the same ideology $m = 0.9$. We want to examine how placing media into a network causes the "average" opinion of the non-media nodes to be closer to the media ideology m than when there are no media nodes; intuitively, this measures the impact of the media on the population of nodes in a network. To establish a baseline against which to measure media impact, we run 50 trials (all figures use 50 trials) with $M = 0$ media nodes. Each of these trials has two sources of randomness: the distribution of initial opinions and the network structure (except when we use a deterministic network structure, such as a complete network). For one trial, we define the baseline media impact to be

$$R_{\text{base}} = \frac{1}{N} \sum_{i=1}^{N} \|x_i^b - m\| \,,$$

where x_i^b is the opinion of account i at convergence. If we average R_{base} over all trials, we obtain the baseline media impact $\overline{R_{\text{base}}}$. We now add $M > 0$ media nodes, with n_M followers each, and run 50 trials for each pair (M, n_M). Each of these trials has three sources of randomness: the distribution of the initial opinions, the network structure (except when we use deterministic networks), and the distribution of media followerships. For each trial at convergence, we define x_i^* to be the opinion of each non-media node. We then calculate

$$R_{\text{inf}} = \frac{1}{N} \sum_{i=1}^{N} \|x_i^* - m\|$$

and average over all trials to get $\overline{R_{\text{inf}}}$. Note that $\overline{R_{\text{inf}}}$ depends on both M and n_M. We define the overall media impact to be

$$R = \frac{\overline{R_{\text{base}}}}{\overline{R_{\text{inf}}}} \,.$$

We can think of R as measuring how much the media has drawn the mean opinion of the population to its own ideology at convergence. First, although it is possible for $\overline{R_{\text{inf}}} > \overline{R_{\text{base}}}$, this would imply that adding media nodes drives the population opinion further away from the media opinion in the opinion space, which seems implausible; throughout all trials, we never observed this result. Hence, we expect that $R \in [1, \infty)$. In cases where the media has done little to impact the population opinion, we expect that $\overline{R_{\text{inf}}}$ is slightly smaller than $\overline{R_{\text{base}}}$, so R should be close to 1. In cases where the media has had a large impact on the population opinion, we intuitively expect the average opinion to be close to the media ideology, so $\overline{R_{\text{inf}}}$ should be small and R should be large. The heatmap (taken from [2]) in Figure 1 shows the media impact R on Facebook friendship networks of Caltech and Reed College [22].

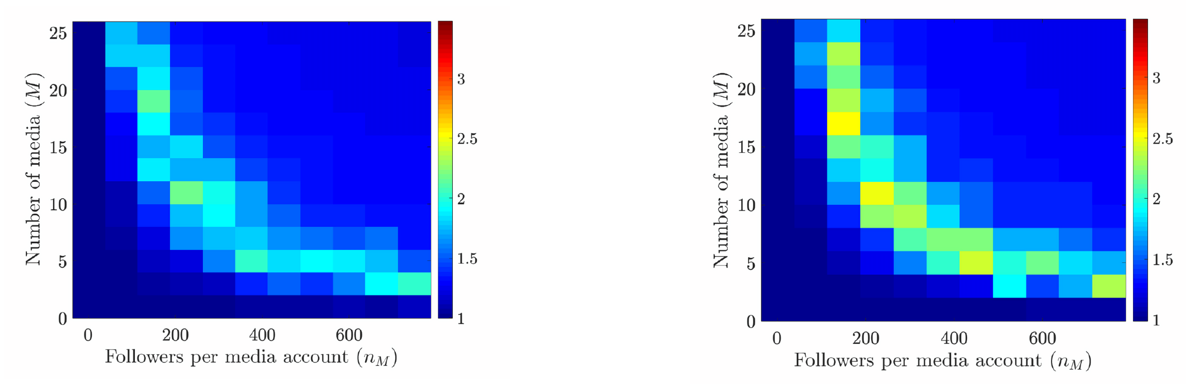

(a) Caltech

(b) Reed College

Figure 1: Heatmaps of media impact R on two college Facebook networks [data taken from [22]]. The vertical axis M is the number of media accounts, and the horizontal axis n_M is the number of followers per media account. Blue represents low acceptance of media ideology, and red represents high acceptance. [Taken from [2]].

Interestingly, past a certain point, increasing M and n_M actually seem to decrease the value of R, which may seem counterintuitive. If one examines a specific trial of the opinions on the Reed College network (data taken from [22]) for $M = 1$, $n_M = 1$ (few media, few followers per media), $M = 11$, $n_M = 225$ (some media, some followers per media), and $M = 21$, $n_M = 675$ (more media, many followers per media), we see that although there are more total media followerships, the mean population opinion is further away from the media ideology, which leads to a decreasing media impact R for increasing M and n_M, as seen in Figure 2.

Finally, in cases where the media has pulled a majority, but not an overwhelming majority, of the population opinion sufficiently close to its own ideology, we see a notable cluster of the population whose opinions are rather far from the media. We believe this notable cluster occurs because the large number of media followerships causes people to more quickly adopt the media's ideology in this case compared to the some-media, some-followerships case. As further evidence, it takes the many-media, many-followerships case far less time than cases with fewer media followerships to reach convergence. Intuitively, the many-media, many-followerships case has "echo chambers" where individuals are too far apart in opinion space to interact. As a result, the media influences a smaller portion of the network.

Unfortunately, R does not capture the qualitative behavior in the many-media, many-followerships case, because the mean opinion is still far from the media ideology. Therefore, we seek to devise another statistic that can help measure the differences between small numbers of media followership and large numbers of media followership. The most obvious quantity to measure is the variance, which should account for how "fractured" a population is. However, there are many other promising statistics and properties to calculate. We now stop following the presentation of [2]; the work in Sections 2–10 is our own.

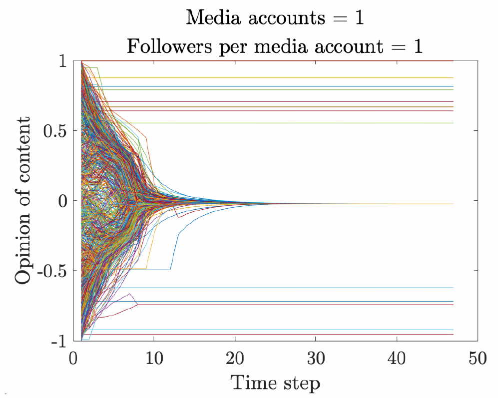

(a) Few media, few followers per media

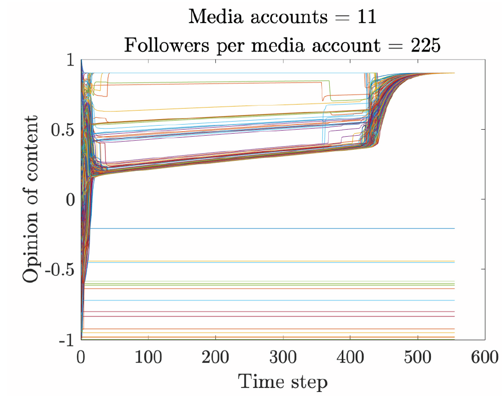

(b) Some media, some followers per media

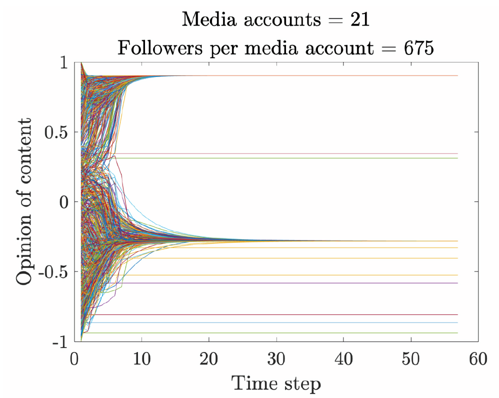

(c) Many media, many followers per media

Figure 2: Time evolution of opinions in the Reed College network for a single trial with various amounts of media followerships. [Taken from [2]].

2 Comparisons of different statistics to measure media impact

From [2], as noted above, R measures how media nodes affect the mean opinion of a population. A large value of R indicates that the media brings the mean opinion of the population close to its own ideology, whereas a small value of R indicates that the media has less of an effect on the mean opinion of the population.

The R measure allows us to compare the influence of the media on the mean opinion in a variety of scenarios. By varying the number of media nodes or the number of followers per media node, we can get a broad picture of the level of *media entrainment*, which measures how much the media has brought the population opinion closer to its ideology. However, R does not clearly distinguish between cases with few media nodes with few followers per node and many media nodes with many followers per node. We explained this issue at the end of Section 1.

Another statistic that we explored is the variance of opinions at convergence:

$$\sigma^2 = \frac{1}{N} \sum_{i=1}^{N} (x_i^* - \mu)^2 \,, \tag{2}$$

where x_i^* is the opinion of individual i at convergence and $\mu = \frac{1}{N} \sum_{i=1}^{N} x_i^*$ is the mean of the individual opinions at convergence. As expected, variance increases with an increasing number of media nodes and an increasing number of followers per media node.

Calculating the variance can indicate how fractured the opinions of the nodes in a network are at convergence, but we did not attempt to gain further information about the number of opinion groups and their distribution.

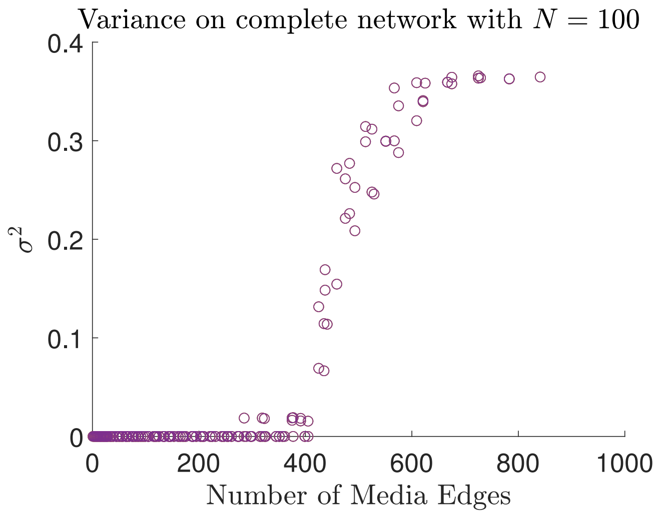

Figure 3: Using Model (1) with a complete network with $N = 100$ nodes, we observe a putative bifurcation in the variance of individual opinions at convergence.

Interestingly, in a complete network with $N = 100$ nodes, when we plot variance at convergence against the total number of media edges (edges $i \to m$ such that individual node i follows media node m, regardless of whether i is receptive to m's content), we find a sudden shift from zero variance to nonzero variance (see Figure 3). This suggests that our model may exhibit a phase transition. Therefore, bifurcation analysis [14] may be able to give further insights into the long-time behavior of the system in (1). See Section 6 for a discussion of bifurcation analysis.

Another statistic that we calculate to measure media influence is

$$\begin{aligned}\chi = \text{ }&\text{proportion of individuals at convergence whose opinion is at most} \\ &\text{the size of the confidence interval } c \text{ from the media ideology.}\end{aligned} \tag{3}$$

The motivation behind the definition of χ is to provide a coarser notion than that of R for measuring the proportion of a population that is influenced by the media. It differs from R in that it does not quantify how close an opinion is from the media ideology; instead, it is concerned only whether an opinion *can* be influenced by the media ideology. So χ captures the difference between many media with many followers and few media with few followers, whereas R does not. However, because it fails to precisely capture how far individual opinions are from the media ideology, it has trouble detecting how small changes in the amount of media followerships affect the distribution of opinions at convergence. We have created heat maps of χ on both complete networks and Watts–Strogatz, but we have not yet fully examined this statistic, because the variance appeared more promising.

Philosophically, there are two different types of media entrainment: influencing individual nodes to share the media ideology or influencing the mean opinion of the population to shift closer to the media ideology. While studying R, we realized that it was better at capturing the latter type of entrainment, rather than the former. To better measure the first type of entrainment with R, we suggest a nonlinear weighting of opinions. For example, in a population with two people and a media ideology of $m = 1$, there should be a quantitative difference between two individuals with opinions $x_1 = x_2 = 0$ and two individuals with opinions $x_1 = 1$ and $x_2 = -1$. Using the original R statistic, there is no difference between these two cases, but we should expect the second case to have a higher entrainment, because individual 2 shares the media ideology. We can then weight the individuals using, for example, a Gaussian distribution with variance s^2 centered at m. In this scenario, the baseline media impact is

$$R^G_{\text{base}} = \frac{1}{N}\sum_{i=1}^{N} \frac{1}{\sqrt{2\pi s^2}} e^{-\frac{\left(x_i^b - m\right)^2}{2s^2}} ,$$

where x_i^b is the opinion of account i at convergence in the case where no media is present. The variance is a parameter that we can adjust. We average R^G_{base} over all trials to calculate the baseline media impact $\overline{R^G_{\text{base}}}$. After incorporating media nodes, we calculate

$$R^G_{\text{inf}} = \frac{1}{N} \sum_{i=1}^{N} \frac{1}{\sqrt{2\pi s^2}} e^{-\frac{\left(x_i^* - m\right)^2}{2s^2}}$$

and average over all trials to get $\overline{R^G_{\text{inf}}}$. We then define the weighted media impact statistic to be

$$R^G = \frac{\overline{R^G_{\text{base}}}}{\overline{R^G_{\text{inf}}}} \,. \tag{4}$$

We have not yet investigated R^G, but we seek to investigate it in the future. Additionally, we can apply a nonlinear weighting of opinions to other statistics to better measure the first type of entrainment.

3 From discrete to continuum in a complete network

Because there appears to be a phase transition in Figure 3 when we measure variance versus the number of media edges in a complete network, we seek to perform a bifurcation analysis of Model (1) in this case. To do this, we formulate a continuum version (continuous in both time and space) of the discrete Model (1). We proceed in four steps to obtain a continuum model:

1. We find an equation for $\{x_i^{\tau,k}\}_{i=1,k=0}^{N,T}$ with fixed time step τ from model (1).
2. We take the limit as $\tau \to 0$ and find an equation for $\{x_i(t)\}_{i=1}^{N}$.
3. We let $N \to \infty$ and derive the corresponding particle-model equation for $x(t)$. (See [3] for a discussion of particle models.) This equation depends on the distribution $\rho(t, x)$ of individuals' opinions and captures the "average" dynamics of the opinion $x(t)$ at time t of a typical individual in a population. In principle, this equation can be derived rigorously by taking a mean-field limit.
4. Using the particle-model equation, we derive the corresponding macroscopic equation for the distribution $\rho(t, x)$ of individual opinions.

3.1 Discrete-in-time particle model

On a complete network with N non-media nodes and M media nodes, the adjacency matrix for our directed network can be written as

$$\mathbf{A} = \left(\begin{array}{c|c} \mathbf{J} & \mathbf{B} \\ \hline 0 & \mathbf{I}_M \end{array}\right),$$

where $\mathbf{A}_{ij} = 1$ if account i follows account j and $\mathbf{A}_{ij} = 0$ otherwise. We also include self-followership. The first $N \times N$ block of $\mathbf{A}$ is $\mathbf{J}$, a matrix of ones (because all non-media nodes follow each other). The first N rows and last M columns give the matrix $\mathbf{B}$, and they encode the non-media nodes that are following media nodes. The last M rows and first N columns give a matrix of zeros (because no media node follows any other nodes). Finally, the last M rows and the last M columns give the identity matrix $\mathbf{I}_M$ (because every media node follows only itself and no other accounts). Note that all of the entries are either 0 or 1.

We define a tuning parameter

$$\alpha = \frac{1}{N} \sum_{j=1}^{M} \sum_{i=1}^{N} \mathbf{B}_{ij} \,. \tag{5}$$

Intuitively we can think of α as measuring the mean number of media followerships per individual.

Consider Model (1): each node x_i (for $i \in \{1, \dots, N+M\}$) at time k follows the update rule

$$x_i^{k+1} = \frac{1}{|I_i|} \left(\sum_{j=1}^{N+M} \mathbf{A}_{ij} x_j^k \eta \left(\left| x_i^k - x_j^k \right| \right) \right),$$

where

$$|I_i| = \sum_{j=1}^{N+M} \mathbf{A}_{ij} \eta \left(\left| x_i^k - x_j^k \right| \right)$$

represents the number of nodes (including media nodes) that i both follows and is receptive to (with respect to content). Recall that for $i \in \{N+1, \dots, N+M\}$, we have that $x_j^k = m$ for all k because media nodes have a fixed ideology. Therefore,

$$x_i^{k+1} = \frac{\sum_{j=1}^{N+M} \mathbf{A}_{ij} x_j^k \eta \left(\left| x_i^k - x_j^k \right| \right)}{\sum_{j=1}^{N+M} \mathbf{A}_{ij} \eta \left(\left| x_i^k - x_j^k \right| \right)} = \frac{\sum_{j=1}^{N} \mathbf{A}_{ij} x_j^k \eta \left(\left| x_i^k - x_j^k \right| \right) + \sum_{j=1}^{M} \mathbf{B}_{ij} m \eta \left(\left| x_i^k - m \right| \right)}{\sum_{j=1}^{N} \mathbf{A}_{ij} \eta \left(\left| x_i^k - x_j^k \right| \right) + \sum_{j=1}^{M} \mathbf{B}_{ij} \eta \left(\left| x_i^k - m \right| \right)}, \tag{6}$$

which is the center of mass of all points to which x_i^k is both receptive and adjacent.

To obtain a continuous-in-time equation, we incorporate a constant scaling factor τ in (6). Let $x_i^{\tau,k}$ represent the opinion of account i at time $k\tau$, where τ is a fixed time step. We define

$$\mathbf{A}^\tau = \left(\begin{array}{c|c} \mathbf{K}^\tau & \tau \mathbf{B} \\ \hline 0 & \mathbf{I}_M \end{array} \right) \text{ where } \mathbf{K}^\tau = \tau \mathbf{J} + (1-\tau) \begin{pmatrix} |I_1| & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & |I_n| \end{pmatrix}.$$

Observe for $i \in \{1, \dots, N\}$ and for all k that

$$\begin{aligned}
|I_i^\tau| &= \sum_{j=1}^{N+M} \mathbf{A}_{ij}^\tau \eta \left(\left| x_i^{\tau,k} - x_j^{\tau,k} \right| \right) \\
&= \sum_{j=1}^{N} \mathbf{K}_{ij}^\tau \eta \left(\left| x_i^{\tau,k} - x_j^{\tau,k} \right| \right) + \sum_{j=1}^{M} \tau \mathbf{B}_{ij} \eta \left(\left| x_i^{\tau,k} - x_j^{\tau,k} \right| \right) \\
&= (1-\tau) |I_i| \eta \left(\left| x_i^{\tau,k} - x_i^{\tau,k} \right| \right) + \sum_{j=1}^{N} \tau \mathbf{J}_{ij} \eta \left(\left| x_i^{\tau,k} - x_j^{\tau,k} \right| \right) + \sum_{j=1}^{M} \tau \mathbf{B}_{ij} \eta \left(\left| x_i^{\tau,k} - x_j^{\tau,k} \right| \right) \\
&= (1-\tau) |I_i| + \tau \sum_{j=1}^{N+M} \mathbf{A}_{ij} \eta \left(\left| x_i^{\tau,k} - x_j^{\tau,k} \right| \right) \\
&= (1-\tau) |I_i| + \tau |I_i| \\
&= |I_i|.
\end{aligned}$$

Also observe that $|I_i^\tau| = 1 = |I_i|$ for all $i \in \{N+1, \dots, N+M\}$ (because these are media nodes). Therefore, $|I_i^\tau| = |I_i|$ for any $i \in \{1, \dots, N+M\}$. We define $\alpha_i = \sum_{j=1}^{M} \mathbf{B}_{ij}$ (the number of media nodes that individual i follows). Note that

$$|I_i| = \sum_{j=1}^{N+M} \mathbf{A}_{ij} \eta \left(|x_i - x_j| \right) = \sum_{j=1}^{N} \eta \left(|x_i - x_j| \right) + \sum_{j=1}^{M} \mathbf{B}_{ij} \eta \left(|x_i - m| \right) = \sum_{j=1}^{N} \eta \left(|x_i - x_j| \right) + \alpha_i \eta \left(|x_i - m| \right).$$

We then propose the new τ-dependent update rule

$$x_i^{\tau,k+1} = \frac{1}{|I_i^\tau|} \sum_{j=1}^{N+M} \mathbf{A}_{ij}^\tau x_j^{\tau,k} \eta \left(\left| x_i^{\tau,k} - x_j^{\tau,k} \right| \right). \tag{7}$$

We see that

$$
\begin{aligned}
x_i^{\tau,k+1} &= \frac{1}{|I_i|}\left(\tau\sum_{j=1}^{N} x_j^{\tau,k}\eta\left(\left|x_i^{\tau,k}-x_j^{\tau,k}\right|\right)+(1-\tau)\left|I_i\right| x_i^{\tau,k}+\tau\sum_{j=1}^{M}\mathbf{B}_{ij} m\eta\left(\left|x_i^{\tau,k}-m_j\right|\right)\right)\\
&=(1-\tau)x_i^{\tau,k}+\frac{\tau}{|I_i|}\left(\sum_{j=1}^{N} x_j^{\tau,k}\eta\left(\left|x_i^{\tau,k}-x_j^{\tau,k}\right|\right)+\alpha_i m\eta\left(\left|x_i^{\tau,k}-m\right|\right)\right),
\end{aligned}
$$

so

$$
\begin{aligned}
\frac{x_i^{\tau,k+1}-x_i^{\tau,k}}{\tau} &= -x_i^{\tau,k}+\frac{1}{|I_i|}\left(\left(\sum_{j=1}^{N} x_j^{\tau,k}\eta\left(\left|x_i^{\tau,k}-x_j^{\tau,k}\right|\right)\right)+\alpha_i m\eta\left(\left|x_i^{\tau,k}-m\right|\right)\right)\\
&=-\frac{1}{|I_i|}\left(\sum_{j=1}^{N}\eta\left(\left|x_i^{\tau,k}-x_j^{\tau,k}\right|\right)+\alpha_i\eta\left(\left|x_i^{\tau,k}-m\right|\right)\right)x_i^{\tau,k}\\
&\quad-\frac{1}{|I_i|}\left(-\left(\sum_{j=1}^{N} x_j^{\tau,k}\eta\left(\left|x_i^{\tau,k}-x_j^{\tau,k}\right|\right)\right)-\alpha_i m\eta\left(\left|x_i^{\tau,k}-m\right|\right)\right)\\
&=-\frac{1}{|I_i|}\left(\sum_{j=1}^{N}\left(x_i^{\tau,k}-x_j^{\tau,k}\right)\eta\left(\left|x_i^{\tau,k}-x_j^{\tau,k}\right|\right)+\alpha_i(x_i^{\tau,k}-m)\eta\left(\left|x_i^{\tau,k}-m\right|\right)\right).
\end{aligned}
$$

Therefore, Model (7) can be written as

$$
\frac{x_i^{\tau,k+1}-x_i^{k,\tau}}{\tau}=-\frac{\frac{1}{N}\sum_{j=1}^{N}\left(x_i^{\tau,k}-x_j^{\tau,k}\right)\eta\left(\left|x_i^{\tau,k}-x_j^{\tau,k}\right|\right)+\frac{\alpha_i}{N}\left(x_i^{\tau,k}-m\right)\eta\left(\left|x_i^{\tau,k}-m\right|\right)}{\frac{1}{N}\sum_{j=1}^{N}\eta\left(\left|x_i^{\tau,k}-x_j^{\tau,k}\right|\right)+\frac{\alpha_i}{N}\eta\left(\left|x_i^{\tau,k}-m\right|\right)}. \tag{8}
$$

Equation (8) reduces to equation (1) for $\tau = 1$.

3.2 Continuous-in-time particle model

It is convenient to express the media nodes and individual nodes as two separate "species" because of their different dynamics. This yields a two-species system; for similar models, see the recent works [4, 12]. We define the piecewise-constant interpolation function

$$
x_i^{\tau}(t) := x_i^{\tau,k} \text{ for } t\in[k\tau,(k+1)\tau), \text{ where } k\in\mathbb{N}.
$$

We also define

$$
x_i(t) := \lim_{\tau\to 0} x_i^{\tau}(t).
$$

From (8), it follows that $x_i(t)$ satisfies the system

$$
\begin{cases}\dot{x}_i = & -\frac{\frac{1}{N}\sum_{j=1}^{N}(x_i-x_j)\eta\left(|x_i-x_j|\right)+\frac{\alpha_i}{N}(x_i-m)\eta\left(|x_i-m|\right)}{\frac{1}{N}\sum_{j=1}^{N}\eta\left(|x_i-x_j|\right)+\frac{\alpha_i}{N}\eta\left(|x_i-m|\right)},\\ \dot{m}_i = & 0,\end{cases} \text{ with initial conditions } \begin{cases}x_i(0) = & x_i^0,\\ m_i(0) = & m,\end{cases} \tag{9}
$$

where x_i represents the ideology of non-media node i and m_j represents the ideology of media node j. Note all x_i and m_j depend on time. However, by the update rule in (7), all of the media nodes maintain the constant ideology m, so $m_j = m$ for all media nodes j at all times. We define the function $\mathcal{K} : \mathbb{R}_{\geq 0} \to \mathbb{R}$ to be such that $\mathcal{K}'(r) = \eta(r)$, and we observe that

$$\begin{aligned}\sum_{j=1}^{N+M} \mathbf{A}_{ij}(x_i - x_j)\eta\left(|x_i - x_j|\right) &= \sum_{j=1}^{N}(x_i - x_j)\eta\left(|x_i - x_j|\right) + \sum_{j=N+1}^{N+M} \mathbf{A}_{ij}(x_i - x_j)\eta\left(|x_i - x_j|\right) \\ &= \sum_{j=1}^{N}(x_i - x_j)\eta\left(|x_i - x_j|\right) + \sum_{j=1}^{M} \mathbf{B}_{ij}(x_i - m_j)\eta\left(|x_i - m_j|\right) \\ &= \sum_{j=1}^{N}(x_i - x_j)\eta\left(|x_i - x_j|\right) + (x_i - m)\eta\left(|x_i - m|\right)\alpha_i \\ &= \sum_{j=1}^{N}(x_i - x_j)\nabla\mathcal{K}\left(|x_i - x_j|\right) + (x_i - m)\nabla\mathcal{K}\left(|x_i - m|\right)\alpha_i \,.\end{aligned}$$

Therefore,

$$\dot{x_i} = \frac{-\left(\frac{1}{N}\sum_{j=1}^{N}(x_i - x_j)\nabla\mathcal{K}\left(|x_i - x_j|\right) + (x_i - m)\nabla\mathcal{K}\left(|x_i - m|\right)\frac{\alpha_i}{N}\right)}{\frac{1}{N}\sum_{j=1}^{N}\eta\left(|x_i - x_j|\right) + \eta\left(|x_i - m|\right)\frac{\alpha_i}{N}} \,. \tag{10}$$

Equation (10) is alternative way to write equation (9). Equation (10) is similar to the discrete aggregation model studied in [3], but (10) is more complicated because of the denominator on the right-hand side.

3.3 Mean-field limit

Let $[-1, 1] = \Omega \subseteq \mathbb{R}^d$ be the opinion space that the individual opinions lie in. Recall that we take $d = 1$ in our investigation. (The work here can be generalized to bounded sets in $\mathbb{R}^d$.) Note that the evolution of opinions in the continuous-in-time particle Model (10) does not depend on the entries of the media followership matrix $\mathbf{B}$, but rather on the number of media each individual is following; i.e., α_i. We thus introduce the function b_N to formulate (10) in such a way that allows us to take the number of non-media nodes to infinity. Let $b_N : \Omega \to \mathbb{R}$ be a function satisfying

$$b_N(x_i) = \sum_{j=1}^{M}\mathbf{B}_{ij} = \alpha_i \,, \quad \text{for} \quad i \in \{1, \ldots, N\}$$

to represent the number of media accounts that non-media node i (with $i \in \{1, \ldots, N\}$) follows. Moreover, we use b_N to encode media followership into our continuum model; we construct b_N ourselves. Assuming that matrix $\mathbf{B}$ is such that the following limit exists, we then define

$$b(x) := \lim_{N\to\infty} \frac{b_N(x)}{N} \,,$$

so $\frac{\alpha_i}{N} \to b(x_i)$ as $N \to \infty$.

If we let $\rho(t, x)$ denote the distribution of individual (non-media) opinions at time t and take the Dirac distribution $\nu(x) = \delta_m(x)$ as the distribution of media opinions, then we may write the mean-field approximation of equation (10) as

$$\dot{x} = \frac{-\int_\Omega (x - y)\eta\left(|x - y|\right)\rho(y)\,dy - \Big((x - m)\eta\left(|x - m|\right)\Big)b(x)}{\int_\Omega \eta\left(|x - y|\right)\rho(y)\,dy + \Big(\eta\left(|x - m|\right)\Big)b(x)} \,, \tag{11}$$

where $x(t) \in \Omega \subseteq \mathbb{R}^d$. Equation (11) expresses the "averaged" particle dynamics and corresponds to the mean-field limit in which we take the number of non-media nodes in the network to infinity, while keeping the number of media nodes fixed. If instead we take the number of media nodes to infinity (in some suitable manner), this is another possible mean-field limit. We discuss this further in Section 10.

3.4 Macroscopic dynamics

From previous work (see [3, 4, 12, 16]), if we know the velocity $\mathbf{v}$ of particles that move as part of a group, the distribution ρ of particles satisfies the continuity equation

$$\partial_t \rho + \mathrm{div}(\mathbf{v}\rho) = 0\,.$$

With the velocity of each particle given by equation (11), the macroscopic density ρ is

$$\partial_t \rho = \mathrm{div}\left(\rho\left(\frac{\int \eta\left(|x-y|\right)(x-y)\rho(t,y)\,dy + (x-m)\eta\left(|x-m|\right)b(x)}{\int \eta\left(|x-y|\right)\rho(t,y)\,dy + \eta\left(|x-m|\right)b(x)}\right)\right)\,, \tag{12}$$

where $b(x)$ represents the influence that media with ideology m exercise on individuals having opinion x.

4 Steady States

We now transition to a discussion of steady states. Often in mathematical modeling problems, the time evolution of a solution is intimately intertwined with its initial conditions. Because the intermediary dynamics are difficult to analyze, we instead examine steady states (the equilibrium behavior of solutions). If we understand the steady states of a system, we simplify a problem's complexity, as we can (hopefully) classify long-term behavior of solutions based on a few initial conditions. In addition, facts about steady states lead to facts about the intermediate behavior of solutions. For example, for a certain configuration of a system, if we know that the variance of opinions in a population at its steady state is 0, then the opinions must converge over time, which by itself is a strong claim. In a perfect world, we would have complete understanding of the steady state of a system given any set of initial conditions. However, this is a very difficult problem, so we turn to probing the steady states with numerical simulation or analytically proving theorems in specific cases. The work in Sections 5–8 is focused on understanding the steady states of Model (12).

From (12), the steady states of the system occur when $\partial_t \rho = 0$. It follows that the steady state ρ_∞ satisfies

$$\int \eta\left(|x-y|\right)(x-y)\rho_\infty(y)\,dy + (x-m)\eta\left(|x-m|\right)b(x) = 0\,. \tag{13}$$

We believe that investigating ρ_∞ for various types of media followerships $b(x)$ will give insight into exactly how the media affects the distribution of opinions at steady-state. Furthermore, we could investigate which choices of $b(x)$ maximize which measures of media impact; this could be helpful for designing advertising strategies (see Section 10). We are currently working on a solver for ρ_∞.

In the case when $b(x) = 0$ (there is no media influence), we give Theorem 1 about the steady states of (12). Note that $b(x) = 0$ corresponds to $\alpha_i = 0$ for all $i \in \{1, \ldots, N\}$; in (9), we then have

$$\dot{x}_i = -\frac{\sum_{j=1}^{N}(x_i - x_j)\eta\left(|x_i - x_j|\right)}{\sum_{j=1}^{N}\eta\left(|x_i - x_j|\right)}$$

for all $i \in \{1, \ldots, N\}$. We use the continuous-in-time particle model in (9) instead of the macroscopic equation in (12) for convenience.

Theorem 1. *At convergence, the opinions of the particle model given by update rule* (9) *without media influence lie in sums of Dirac distributions centered at points at least c apart; i.e. they take the form*

$$\rho_{conv}(z) = \sum_{i=1}^{K}\beta_i\delta_{y_i}(z),$$

for some integer K and selection of points $\{y_1, \ldots, y_K\}$, where each δ_{y_i} is centered at least c apart and the coefficients β_i satisfy $\sum_{i=1}^{K}\beta_i = 1$.

Proof. There must exist some x such that $\rho_{\text{conv}}(x) > 0$. Choose i such that $\rho_{\text{conv}}(x_i) > 0$ and x_i is maximized. We claim that there is no other opinion $x_j \neq x_i$ such that $\eta\left(|x_i - x_j|\right) = 1$. Suppose, by contradiction, that there exists an x_j such that $|x_i - x_j| = x_i - x_j < c$. Because x_i is the maximum point at which ρ_{conv} is nonzero, we know that $\rho_{\text{conv}}(x) = 0$ for any $x > x_i$. We conclude there exists an x_j such that

$$(x_i - x_j)\eta\left(|x_i - x_j|\right) = x_i - x_j > 0.$$

If no such x_j existed, then all the opinions would lie in a Dirac centered at x_i, so assume this is not the case. Since $0 \leq \eta\left(|x_i - x|\right) \leq 1$ for all x, we know that $\sum_{j=1}^{N} \eta\left(|x_i - x_j|\right) > 0$. Therefore,

$$\dot{x}_i = -\frac{\sum_{j=1}^{N}(x_i - x_j)\eta\left(|x_i - x_j|\right)}{\sum_{j=1}^{N}\eta\left(|x_i - x_j|\right)} < 0,$$

since the denominator is positive and the numerator has only nonnegative summands (with at least one positive summand). However, because we are at steady state, we must have $\dot{x}_i = 0$. Therefore, we are not in a steady state, and there cannot possibly exist such an x_j.

Therefore, x_i must lie in a Dirac distribution with center at least c away from any other opinion. Consequently, the opinions centered at x_i do not interact with the other opinions. If we remove the Dirac distribution at x_i and renormalize, ρ_{conv} will therefore still be a steady state. If we repeat the above argument for the normalized ρ_{conv}, we then see that the new maximum opinion x_i' that we pick after the normalization is at least c below the old x_i. Therefore, this process can only repeat finitely many times. Repeating this argument, we see that all opinions must lie in Dirac distributions with centers at least c apart from each other. □

5 Variance for the discrete-in-time model in (7) versus the continuous-in-time model in (9)

We can calculate the variance (one of the media impact measures mentioned in Section 2) of non-media opinions at steady states to understand how fractured a population's opinions are. As mentioned in Section 4, there are several initial conditions that affect the variance. We consider some of these in Section 5. We would like to numerically investigate the behavior of the discrete-in-time model in (7) as the time step τ decreases (approaching the continuous-in-time model in (9)). The behavior of solutions to both models are dependent on the choice of initial condition for the opinions and the way in which media followership is increased. In Section 5, we concentrate on the former. We address the latter in Section 6.

After developing Model (7) with an explicit fixed time step τ, we ran simulations with $\tau = 0.1$ on a complete network with $N = 100$ nodes. We choose the media followerships by constructing a $\mathbf{B}$ matrix with nonzero entries appearing uniformly at random. We chose the initial population opinions in two ways:

1. **Initial Condition 1: linear spacing.** We linearly space N opinions from the left endpoint of the opinion space to the right endpoint of the opinion space, inclusive of the end points. This initial condition is deterministic.

2. **Initial Condition 2: uniformly at random.** We draw N opinions independently at random from a uniform distribution on the opinion space. This initial condition is non-deterministic.

Using a deterministic linear spacing of initial opinions, we obtained the plot in Figure 4(a). We examined the same parameters in the continuum Model (9), and we obtained a similar result shown in Figure 4(b). We also observe a similarity in shape between $\tau = .1$ and the continuum model when the initial opinions are drawn uniformly at random; see Figure 5.

As we have observed numerically, the discrete model converges to the continuous-in-time model; therefore, when we study bifurcations in Section 6, we use (9) for numerics because it runs faster than (7). Additionally, there is a difference between choosing the initial opinions uniformly at random and linearly spacing them; the former yields no observable bifurcation in any of our trials, whereas the latter yields an observable bifurcation under certain conditions. Given enough trials, we expect that the former converges to the latter, but we have not yet fully investigated this.

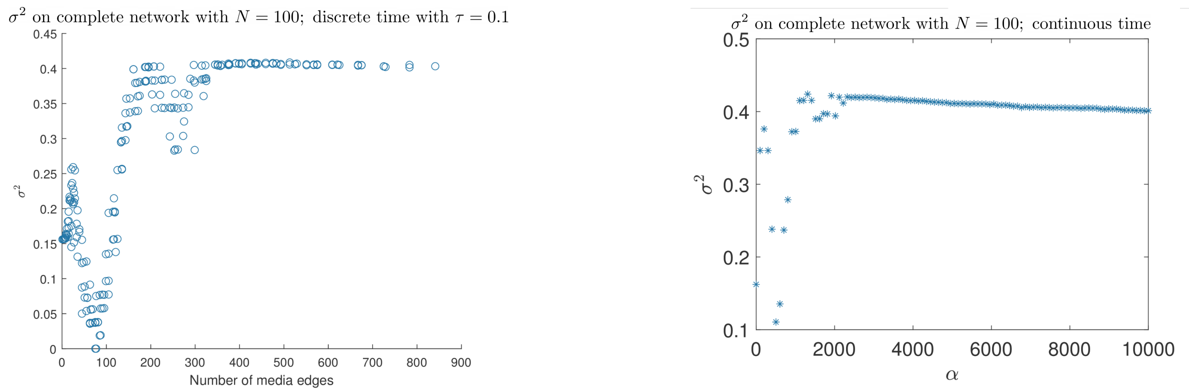

(a) Discrete-in-time Model (7) using a time step $\tau = 0.1$ with initial condition 1

(b) Continuous-in-time Model (9) with initial condition 1

Figure 4: We plot the variance of the individual opinions at convergence for the two models on a complete network with $N = 100$ nodes using initial condition 1. We choose the entries of the $\mathbf{B}$ matrix to be 1 uniformly at random. We define the variance σ^2 in (2) and the number of media edges α in (5).

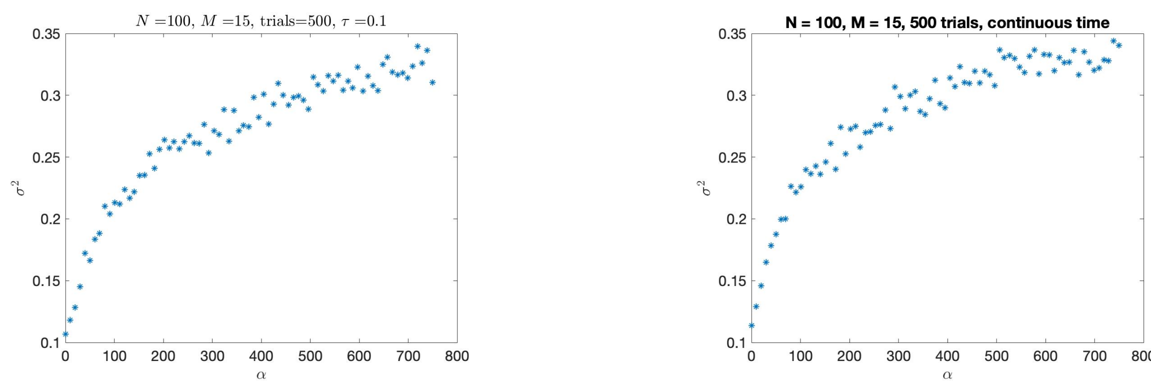

(a) Discrete-in-time Model (7) using time step $\tau = 0.1$ with initial condition 2

(b) Continuous-in-time Model (9) with initial condition 2

Figure 5: We plot the variance of the individual opinions at convergence for Model (7) and Model (9) on a complete network with $N = 100$ nodes using initial condition 2. We choose the entries of the $\mathbf{B}$ matrix to be 1 uniformly at random. We define the variance σ^2 in (2) and the number of media edges α in (5).

6 Numerical bifurcation analysis

Models often have parameters that appear in their governing equations. Adjusting these parameters can change the qualitative behavior of these models, including at steady state. When the qualitative dynamics of a model change with the value of a parameter, one says that there is a *bifurcation* [14]. From our numerical computations on the discrete model in (7), we believe that a bifurcation occurs in the continuum model in (12), at least for certain types of media followership.

To get started on bifurcation analysis, we reviewed [21] to see if the analytical methods that have been developed by Kuramoto and others who have studied Kuramoto-type models are relevant to our analysis. However, his methods are not directly applicable, because Model (9) and Model (12) do not have a simple aggregation equation.

A rigorous proof of a bifurcation is often mathematically involved. For example, J. D. Crawford proved the existence of a bifurcation in the Kuramoto model in the papers [5–7], which contain rather intricate analysis. To prove that a bifurcation exists, we must first show that (12) is the continuum limit of (9). A rigorous proof of this mean-field limit is left as future work. We then must analytically show a qualitative change in the state (a nonzero variance) using (12) as we adjust some tuning parameter [14]. We discuss our specific choice of tuning parameter below. We use linearly-spaced initial opinions for all work in Section 6.

We tested four possible strategies for adjusting the tuning parameter (i.e., the amount of media followership) of our bifurcation. These four strategies are different methods of increasing media followership across the network; they correspond to four different ways of filling the **B** matrix with ones. (A fifth way of filling the **B** matrix would be choosing entries to be 1 uniformly at random; this produced Figures 4 and 5.) The strategies are

1. **Strategy (1): filling in the B matrix by rows.** We fill the **B** matrix by rows from top to bottom. We fill entries in a row one at a time; after we fill a row, we start filling the highest unfilled row. For linearly-spaced initial opinions, this strategy yields $\alpha_i = M$ for individuals close to the -1 opinion and $\alpha_i = 0$ for individuals close to the 1 opinion.

2. **Strategy (2): filling in the B matrix by columns.** We fill the **B** matrix by columns from left to right. We fill entries in a column one at a time; after we fill a column, we start filling the rightmost unfilled column. This strategy yields roughly equal α_i for all individuals.

3. **Strategy (3): filling in the B matrix by randomly selecting rows.** We select a permutation of the rows of **B** uniformly at random and fill the rows in permutation order.

4. **Strategy (4): filling in the B matrix by diagonals.** We fill the **B** matrix by diagonals, starting from the upper-right diagonal and stopping at the main diagonal. This strategy roughly yields $\alpha_1 = k, \alpha_2 = k-1, \ldots, \alpha_k = 0, \alpha_{k+1} = 0, \ldots$ where individual 1 is closest to the -1 opinion.

We observed a bifurcation when employing Strategies (1) and (2). We did all our numerics using the continuous-in-time model in (9) with linearly-spaced opinions as the initial condition, under the assumption that simulations with randomly-spaced initial opinions converge to their counterparts with linearly-spaced initial opinions after sufficiently many trials. See Section 5 for a discussion of our decision to use linearly-spaced initial opinions.

6.1 Filling in the B matrix by rows

For our simulations with $N = 100$ nodes, we observe a bifurcation (see Figure 6):

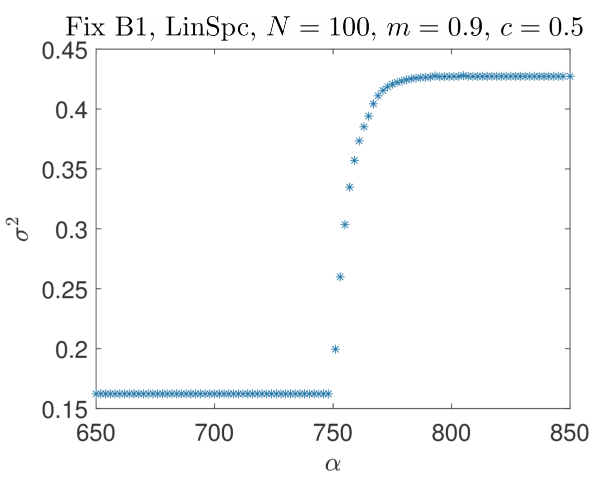

Figure 6: Using Model (9), we observe a putative bifurcation in the variance of individual opinions at convergence for a complete network with $N = 100$ nodes. We choose the entries of **B** using Strategy (1). We define the variance σ^2 in (2) and the number of media edges α in (5).

Intuitively, increasing α using Strategy (1) of filling the $\mathbf{B}$ matrix corresponds to gradually increasing the media influence starting from opinions near -1 to opinions near 1. However, this scenario is rather unrealistic (we do not believe that there is empirical evidence of the media only targeting people whose opinions are close to -1), so we want to examine other ways of filling $\mathbf{B}$ that yield a bifurcation.

6.2 Filling in the B matrix by columns

For $N = 100$ nodes, we did not observe any clear bifurcation. We then tried a variant of Strategy (2) in which we fill in only the center third of the $\mathbf{B}$ matrix. In this strategy, the media only influences these more "moderate" opinions. We do not observe a bifurcation for networks with more than $N = 300$ nodes.

We also examined filling in the middle half of the $\mathbf{B}$ matrix, instead of the middle third. This leads to a clear bifurcation; the plots in Figure 7 show the bifurcation for $N = 300$, $N = 500$, $N = 750$, and $N = 1000$ nodes.

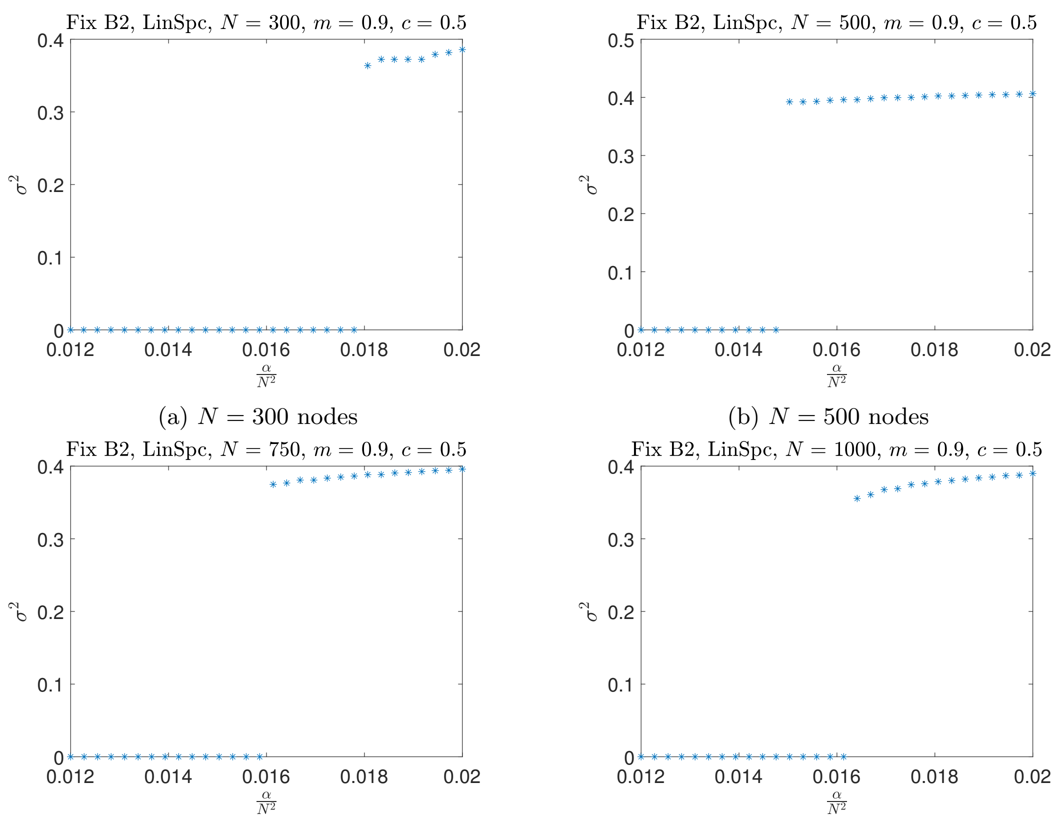

(a) $N = 300$ nodes

(b) $N = 500$ nodes

(c) $N = 750$ nodes

(d) $N = 1000$ nodes

Figure 7: We plot the variance of individuals opinions at convergence. Using Model (9) and a modified version of Strategy (2) (only filling in the middle half of $\mathbf{B}$) to fill the $\mathbf{B}$ matrix, we observe a bifurcation for multiple values of N whose critical point α_c grows roughly with N^2. We define the variance σ^2 in (2) and the number of media edges α in (5).

We also list the critical point α_c for various N in Table 1. The critical point α_c appears to grow approximately with N^2.

N	300	500	750	1000	5000
α_c	1609	3718	8898	16316	439901
$\frac{\alpha_c}{N^2}$	.017878	.014872	.015819	.016316	.017596

Table 1: A table of network sizes N and their respective critical points α_c (rounded to six decimal places) for media ideology $m = 0.9$. We use Model (9) and fill the entries of **B** using the modified version of Strategy (2), where we only fill in the middle half of **B**. Here, α_c, the bifurcation point, is the first value of α such that $\sigma^2 > 0.10$. Interestingly, α_c/N^2 is roughly constant as N increases.

We then tested filling in the lower half of the **B** matrix (corresponding roughly to opinions from 0 to 1) and observed no bifurcation. Filling in the upper half of the **B** matrix (corresponding roughly to opinions from -1 to 0) also has no bifurcation.

We also tried filling in lower 3/4 of the **B** matrix (corresponding roughly to opinions from -0.5 to 1); in this case, we observed a bifurcation. We next filled in the upper 3/4 of the **B** matrix (corresponding roughly to opinions from -1 to 0.5) and did not observe a bifurcation.

Notably, the continuum limit of Strategy (2) for $b(x)$ is a box function. We chose this strategy because we believed that influencing these moderate opinions would lead to a bifurcation.

6.3 Filling in the B matrix by randomly selecting rows

We do not observe any bifurcation for $N = 100$ or $N = 300$ nodes. Even when we restrict the randomly chosen rows to the middle half of the rows of the **B** matrix, we did not observe a bifurcation. We chose this strategy because we believed that increasing the number of people following media nodes (i.e., increasing the number of nonzero rows of the **B** matrix) would lead to a bifurcation. We believe that the lack of an observed bifurcation for $N = 300$ nodes is because no such bifurcation exists (as was seen in Strategy (1)). In the case of $N = 100$ nodes, we suspect that the observed bifurcation is sensitive to which individuals are influenced by the media; however, we are interested in finding a concrete explanation for the lack of an observed bifurcation in this case.

6.4 Filling in the B matrix by diagonals

We see what appears to be a bifurcation for $N = 100$ nodes, but we did not observe a bifurcation for $N = 300$ nodes, which leads us to believe that the putative bifurcation in $N = 100$ nodes was an artifact of the finite size of our networks. We chose this strategy to test whether the distribution of media followerships among opinions was important in the existence of a bifurcation.

7 Numerical effects of varying c and m on a complete network

7.1 Different consensus groups based on c

To gain a sense of the qualitative behavior of Model (9) with different levels of media influence, we plot the number of opinion groups at $t = 1000$ (note that we round each opinion to four decimal places) for $N = 300$ nodes with $\alpha = 0$, $\alpha = 600$, $\alpha = 1200$, and $\alpha = 1800$. Intuitively, when $\alpha = 0$, we expect the population to converge to consensus. Therefore, Figure 8 gives an idea of the minimum c that we need to reach consensus. Interestingly, the smallest c for $N = 300$ nodes is close to 0.5, which is exactly the choice of Brooks and Porter [2].

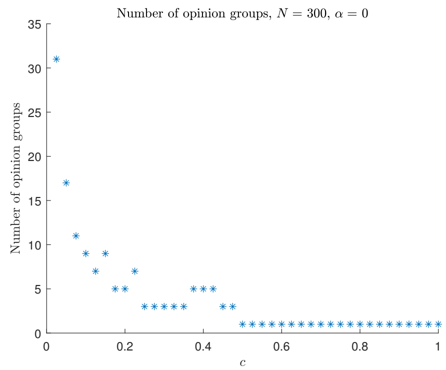

(a) Opinion groups for $N = 300$ nodes and $\alpha = 0$

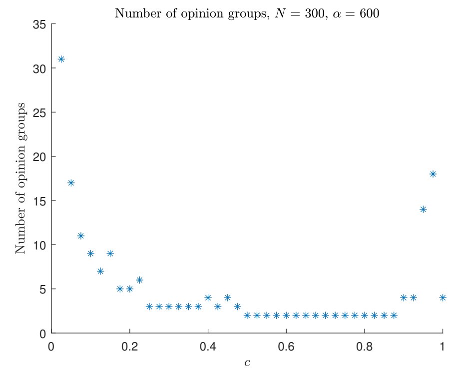

(b) Opinion groups for $N = 300$ nodes and $\alpha = 600$

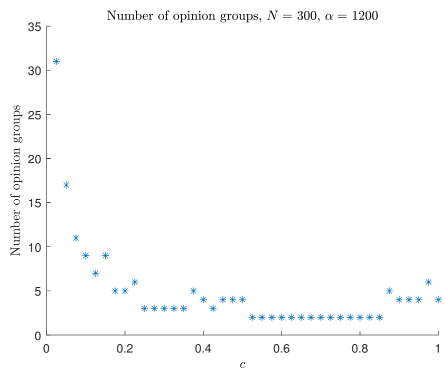

(c) Opinion groups for $N = 300$ nodes and $\alpha = 1200$

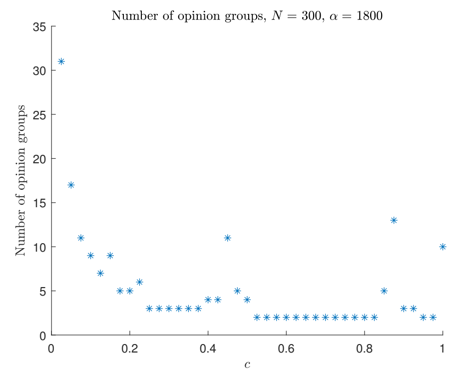

(d) Opinion groups for $N = 300$ nodes and $\alpha = 1800$

Figure 8: For progressively larger α, the number of opinion groups in a complete network with $N = 300$ nodes for most c at $t = 1000$ appears to become larger, which suggests that the population is growing increasingly polarized.

7.2 Varying media ideology

We plot our results of simulations when we employ Strategy (2) with $N = 500$ nodes and media ideologies $m = 0.8$, $m = 0.85$, $m = 0.95$, and $m = 1.0$ (see Figure 9) to see if the bifurcation point itself changes. It appears that the bifurcation is sensitive to the value of m (see Table 2). This may suggest that for any given N, there exists a value of m that minimizes α_c.

m	.95	.90	.85	.80	.75	.70	.65	.60	.55	.50
α_c	5249	3718	3707	4378	5693	6746	9676	9499	-	-
$\frac{\alpha_c}{N^2}$	.020996	.014872	.014828	.017512	.022772	.026984	.038704	.037996	-	-

Table 2: A table of media ideologies m and their respective critical points α_c (rounded to six decimal places) for complete networks with $N = 500$ nodes. We use Model (9) and fill the entries of $\mathbf{B}$ using the modified version of Strategy (2), where we only fill in the middle half of $\mathbf{B}$. Here, α_c is the first value of α such that $\sigma^2 > 0.10$. An empty entry means we do not observe a bifurcation for that value of m.

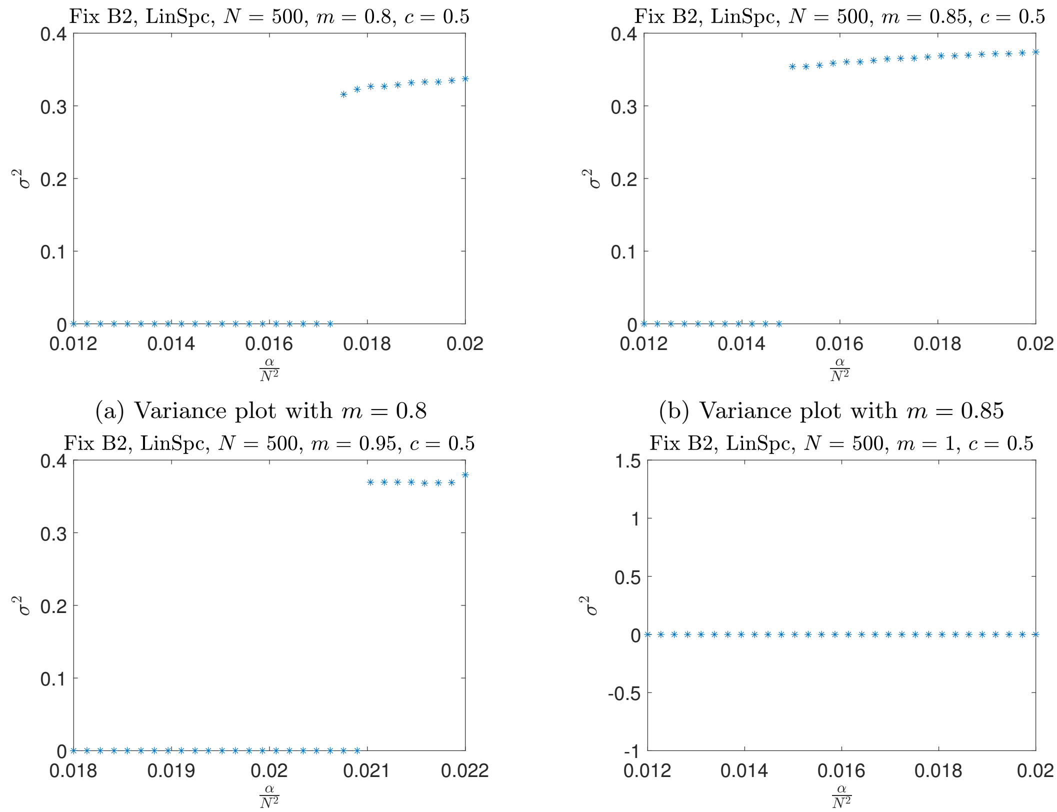

(a) Variance plot with $m = 0.8$

(b) Variance plot with $m = 0.85$

(c) Variance plot with $m = 0.95$

(d) Variance plot with $m = 1.0$

Figure 9: For progressively larger m, the critical point α_c of the bifurcation appears to shift to the left. It seems that after m increases past a threshold in a complete network with $N = 500$ nodes, we do not observe a bifurcation.

8 Adjustments to the model in (1)

8.1 Using a symmetric update rule

Here we explore how a complete network without media nodes should behave at convergence. Intuitively, we expect $\bar{x}$, the mean of the non-media opinions, to be conserved over time. We also expect the time derivative of $\sigma^2(x)$, the second moment of the opinions, to be monotonically non-increasing as individuals evolve towards consensus within groups to which they are receptive. Unfortunately, we were not able to show these results using Model (9). Demonstrating these results is challenging because the denominator is not symmetric with respect to our summation indices, so we will now adjust the model to make the mathematical analysis more tractable. We start by defining

$$d_i := \sum_{j=1}^{N} \eta\left(|x_i - x_j|\right) \tag{14}$$

to represent the "out-degree" of any node. In Model (9), in the absence of media, the time derivative of the mean $\bar{x}(t) = \frac{1}{N}\sum_{i=1}^{N} x_i(t)$ is

$$\frac{d}{dt}\left(\bar{x}(t)\right) = \frac{1}{N}\sum_{i=1}^{N}\dot{x}_i(t) = -\frac{1}{N}\sum_{i=1}^{N}\sum_{j=1}^{N}\frac{(x_i - x_j)\eta\left(|x_i - x_j|\right)}{d_i} = \frac{1}{2N}\sum_{i,j}\left(\frac{1}{d_i} - \frac{1}{d_j}\right)(x_i - x_j)\eta\left(|x_i - x_j|\right) ,$$

and it is not clear whether this vanishes. Also note that this means that once the distribution of opinions has reached a situation where opinions are clustered in groups of diameter at most c, then the mean remains constant. If we define the new update rule

$$\dot{x}_i(t) = -N \sum_{j=1}^{N} \frac{(x_i - x_j)\eta\left(|x_i - x_j|\right)}{d_i d_j}\,, \tag{15}$$

we obtain

$$\begin{aligned}
\frac{d}{dt}(\bar{x}_i(t)) &= -\sum_{i=1}^{N}\sum_{j=1}^{N} \frac{(x_i - x_j)\eta\left(|x_i - x_j|\right)}{d_i d_j} \\
&= -\sum_{i=1}^{N}\sum_{j=1}^{N} \frac{x_i\eta\left(|x_i - x_j|\right)}{d_i d_j} + \sum_{i=1}^{N}\sum_{j=1}^{N} \frac{x_j\eta\left(|x_i - x_j|\right)}{d_i d_j} \\
&= 0\,,
\end{aligned}$$

where the last equality follows from exchanging the summation indices in the second term. The time derivative of the second moment $\sigma^2(t) = \frac{1}{N}\sum_{i=1}^{N}(x_i(t))^2$ is also not easy to analyze in Model (9). However, using the update rule (15), the time derivative of the second moment is

$$\begin{aligned}
\frac{d}{dt}\left(\sigma^2(t)\right) &= \frac{2}{N}\sum_{i=1}^{N} x_i(t)\dot{x}_i(t) \\
&= -2\sum_{i=1}^{N}\sum_{j=1}^{N} \frac{x_i(x_i - x_j)\eta\left(|x_i - x_j|\right)}{d_i d_j} \\
&= -2\left(\sum_{i=1}^{N}\sum_{j=1}^{N} \frac{x_i^2\eta\left(|x_i - x_j|\right)}{d_i d_j} - \sum_{i=1}^{N}\sum_{j=1}^{N} \frac{x_i x_j\eta\left(|x_i - x_j|\right)}{d_i d_j}\right) \\
&= -2\left(\frac{1}{2}\sum_{i=1}^{N}\sum_{j=1}^{N} \frac{x_i^2\eta\left(|x_i - x_j|\right)}{d_i d_j} - \sum_{i=1}^{N}\sum_{j=1}^{N} \frac{x_i x_j\eta\left(|x_i - x_j|\right)}{d_i d_j} + \frac{1}{2}\sum_{i=1}^{N}\sum_{j=1}^{N} \frac{x_j^2\eta\left(|x_i - x_j|\right)}{d_i d_j}\right) \\
&= -\left(\sum_{i=1}^{N}\sum_{j=1}^{N} \frac{(x_i - x_j)^2\eta\left(|x_i - x_j|\right)}{d_i d_j}\right) \\
&\leq 0\,,
\end{aligned}$$

because all terms in the parentheses are always non-negative. This shows that the time derivative of the second moment is monotonically non-increasing. Observe that it is equal to 0 only when either $\eta(|x_i - x_j|) = 0$ or $x_i = x_j$ for all pairs i, j; this implies that all of the opinions concentrate on points that are at least c apart. In that case, we also know that $\dot{x}_i = 0$ for all i, so the dynamics have converged.

Model (9) without media influence can be interpreted as a variation of opinion dynamics driven by a graph Laplacian; we will make this idea precise below. Consider the normalized graph Laplacian

$$L = D^{-1/2}(D - W)D^{-1/2}\,,$$

where $W_{ij} \coloneqq \eta\left(|x_i - x_j|\right)$ and $D \coloneqq \mathrm{diag}(d_i)$, where $d_i = \sum_j W_{ij}$ [23]. Let $\mathbf{x} = (x_1, \ldots, x_N)$ be a vector of all the opinions. Then consider the update rule

$$\begin{aligned}
\dot{\mathbf{x}} &= -L\mathbf{x} \\
&= -\begin{bmatrix} d_1^{-1/2} & \cdots & 0 & \cdots & 0 \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ 0 & \cdots & d_i^{-1/2} & \cdots & 0 \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ 0 & \cdots & 0 & \cdots & d_N^{-1/2} \end{bmatrix} \begin{bmatrix} d_1 - 1 & \cdots & -\eta\left(|x_1 - x_i|\right) & \cdots & -\eta\left(|x_1 - x_N|\right) \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ -\eta\left(|x_i - x_1|\right) & \cdots & d_i - 1 & \cdots & -\eta\left(|x_i - x_N|\right) \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ -\eta\left(|x_N - x_1|\right) & \cdots & -\eta\left(|x_N - x_i|\right) & \cdots & d_N - 1 \end{bmatrix} \\
&\quad \begin{bmatrix} d_1^{-1/2} & \cdots & 0 & \cdots & 0 \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ 0 & \cdots & d_i^{-1/2} & \cdots & 0 \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ 0 & \cdots & 0 & \cdots & d_N^{-1/2} \end{bmatrix} \begin{bmatrix} x_1 \\ \vdots \\ x_i \\ \vdots \\ x_N \end{bmatrix} \\
&= -\begin{bmatrix} d_1^{1/2} - d_1^{-1/2} & \cdots & -d_1^{-1/2}\eta\left(|x_1 - x_i|\right) & \cdots & -d_1^{-1/2}\eta\left(|x_1 - x_N|\right) \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ -d_i^{-1/2}\eta\left(|x_i - x_1|\right) & \cdots & d_i^{1/2} - d_i^{-1/2} & \cdots & -d_i^{-1/2}\eta\left(|x_i - x_N|\right) \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ -d_N^{-1/2}\eta\left(|x_N - x_1|\right) & \cdots & -d_N^{-1/2}\eta\left(|x_N - x_i|\right) & \cdots & d_N^{1/2} - d_N^{-1/2} \end{bmatrix} \begin{bmatrix} d_1^{-1/2} x_1 \\ \vdots \\ d_i^{-1/2} x_i \\ \vdots \\ d_N^{-1/2} x_N \end{bmatrix} \\
&= \begin{bmatrix} -x_1 + \sum_{j=1}^N \frac{\eta\left(|x_1 - x_j|\right) x_j}{d_1^{1/2} d_j^{1/2}} \\ \vdots \\ -x_i + \sum_{j=1}^N \frac{\eta\left(|x_i - x_j|\right) x_j}{d_i^{1/2} d_j^{1/2}} \\ \vdots \\ -x_N + \sum_{j=1}^N \frac{\eta\left(|x_N - x_j|\right) x_j}{d_N^{1/2} d_j^{1/2}} \end{bmatrix}
\end{aligned}$$

Therefore, we introduce the update rule

$$\dot{x}_i = -x_i + \sum_j \frac{\eta\left(|x_i - x_j|\right) x_j}{d_i^{1/2} d_j^{1/2}} = -\sum_j \frac{\eta\left(|x_i - x_j|\right)}{d_i^{1/2}} \left(\frac{x_i}{d_i^{1/2}} - \frac{x_j}{d_j^{1/2}} \right), \tag{16}$$

where d_i is as defined in (14). As shown, the update rule is motivated by the graph Laplacian. We prove Theorem 2 for the steady states of (16).

Theorem 2. *At convergence, opinions in the particle model given by update rule* (16) *lie in sums of Dirac distributions centered at points at least c apart; i.e. they take the form*

$$\rho_{conv}(z) = \sum_{i=1}^K \beta_i \delta_{y_i}(z),$$

for some integer K and selection of points $\{y_1, \ldots, y_K\}$, where each δ_{y_i} is centered at least c apart and the coefficients β_i satisfy $\sum_{i=1}^K \beta_i = 1$.

Proof. Let $\psi(x)$ be the proportion of the nodes with opinion at most c from opinion x. Therefore, if x_i is the opinion of node i, then

$$d_i = \sum_j \eta\left(|x_i - x_j|\right) = N\psi(x_i)\,,$$

where N is the total number of nodes. Let ρ_{conv} be the distribution of nodes over the opinion space $[-1, 1]$ at steady state. Normalize ρ_{conv} such that $\int_{-1}^{1} \rho_{\text{conv}}(z)\, dz = 1$. Note that

$$\psi(x) = \int_{x-c}^{x+c} \rho_{\text{conv}}(z)\, dz, \text{ so } d_i = N\psi(x_i) = N \int_{x_i-c}^{x_i+c} \rho_{\text{conv}}(z)\, dz$$

and that $\rho_{\text{conv}}(x) = 0$ if $|x| > 1$.

There must exist some x such that $\rho_{\text{conv}}(x) > 0$. Choose i such that $\rho_{\text{conv}}(x_i) > 0$ and $|x_i|$ is maximized. If $|x_i| = 0$, then all opinions are centered on a Dirac distribution at $x = 0$; this satisfies the claim. Therefore, assume that this is not the case. Because the update rule is symmetric with respect to $x = 0$, we may assume that $x_i > 0$ without loss of generality.

We claim that there is no other opinion $x_j \neq x_i$ such that $\eta\left(|x_i - x_j|\right) = 1$. Suppose, by contradiction, that there exists an x_j such that $|x_i - x_j| = x_i - x_j < c$. Because x_i is the maximum point at which ρ_{conv} is nonzero, we know that $\rho_{\text{conv}}(x) = 0$ for any $x > x_i$. Therefore,

$$d_i = N \int_{x_i-c}^{x_i+c} \rho_{\text{conv}}(z)\, dz = N \int_{x_i-c}^{x_i} \rho_{\text{conv}}(z)\, dz \leq N \int_{x_j-c}^{x_i} \rho_{\text{conv}}(z)\, dz = N \int_{x_j-c}^{x_j+c} \rho_{\text{conv}}(z)\, dz = d_j\,. \quad (17)$$

Therefore, because $x_i > x_j$, it cannot be the case that $x_i/d_i = x_j/d_j$. Therefore, we are not in a steady state, and there cannot possibly exist such an x_j.

Therefore, x_i must lie in a Dirac distribution with center at least c away from any other opinion. Consequently, the opinions centered at x_i do not interact with the other opinions. If we remove the Dirac distribution at x_i and renormalize, ρ_{conv} will therefore still be a steady state. If we repeat the above argument for the normalized ρ_{conv}, we then see that the new opinion x_i' that we pick after the normalization is at least c below the old x_i. Therefore, this process can only repeat finitely many times. Repeating this argument, we see that all opinions must lie in Dirac distributions with centers at least c apart from each other. □

Unfortunately, although we can add a symmetry term d_j to the continuum Model (9) to make the analysis easier, we do not yet know where to add the symmetry term d_j to the discrete Model (7). More work needs to be done to determine where we can add d_j. From our current perspective, although they are mathematically interesting, the models (15) and (16) are conceptually different from our original model. We leave them for further research efforts.

9 Conclusions

We began this SURF project off of the work done by Brooks and Porter in [2]. We wanted to better understand how media impacts opinion in society. To make this problem approachable mathematically, we made several simplifications, such as: representing the population with a network, modeling social interactions with an opinion model, and devising measures to gauge "media impact".

We devised several measures of media impact, such as σ^2 in (2), χ in (3), and R^G in (4). We focused on the variance of opinions at convergence and how it behaves with various initial conditions, network structure, and amount of media influence. In one of these trials, we observed a bifurcation in Model (1), leading us to develop a mean-field limit (see Equation (12)). We proved Theorem 1 regarding the steady states of Model (9). After running numerical simulations with the continuous-in-time Model (9), we realized the bifurcation was sensitive to both initial conditions and how the tuning parameter, the amount of media influence, is increased. We also experimented with other types of models, one of which resulted from exploring the graph Laplacian, and proved Theorem 2 regarding steady states of Model (16). At this point, there are several avenues for further research.

10 Future plans

Since we did not expect to observe a bifurcation in Model (1), we now want to rigorously prove the existence of a bifurcation for Model (12) and understand how the long term behavior of solutions depends on the tuning parameter α, the amount of media influence. Additionally, we want to explore the behavior of the solution ρ to the macroscopic Model (12) with numerics to gain a better understanding of the observed bifurcation using Strategy (2). As a first step, we plan to implement a numerical scheme to simulate ρ_∞ from (13). One possible idea to explore is if the bifurcation holds when we decrease c as some function of N as N increases. This represents holding constant the number of other individuals with whom any one individual interacts. Such a continuum model is entirely different from Model (12) in Section 3. One argument in favor of this continuum limit is the principle of cognitive capacity [11], which implies that individuals realistically only interact with a finite number of other individuals at any given time. Consequently, keeping the number of individuals in a characteristic neighborhood (the confidence interval) constant is a reasonable choice to explore. We also want to examine the possibility that for any given complete network size N, there exists a media ideology m that minimizes the critical point α_c. Table 2 supports this possibility.

We seek to use spectral clustering techniques to determine opinion groups in a population. Specifically, we want to find a method to separate opinions into clusters and determine the number, size, and spread of opinion clusters at convergence. See [23] for more details on spectral clustering techniques. We initially planned on studying the graph Laplacian, leading to development of Model (16), which is related to the graph Laplacian. Even though the model variants in Section 8 are not directly related to Model (1) (the one we focused on), we are interested in exploring their connection to the graph Laplacian. Because the graph Laplacian is well studied in other areas, there are known properties that we can use to analyze the dynamics. Also, because there are no laws that govern social dynamics (to our knowledge), we might be able to make more analytical insights with a model that is more mathematically tractable (for example: choosing a model whose moments are explicit in the absence of media). Finally, if we understand the dynamics of a similar model, we might be able to draw similar insights with Model (1).

We are also interested in further exploring the susceptibility of a population on a given network to media influence. Instead of using the media impact statistic R, one can calculate χ in (3), the fraction of nodes in a network whose opinion is within the confidence interval of the media ideology. We think that this intuitively captures how receptive a population is to the media, but it is likely too coarse of a measure. It does not distinguish between many nodes that are clustered very closely to the media ideology versus pockets of individuals who are all somewhat close to the media ideology (but not quite at it). We want to come up with another statistic to help quantify this idea.

We believe that R^G will measure media entrainment with a finer distinction than R by removing cases where an opinion near -1 and an opinion near 1 "cancel" in the mean. We plan to run numerics to gain intuition on the behavior of R^G under different initial conditions and network structures.

Moreover, we are looking into which sorts of algorithms (ways to add media influence into the opinion dynamics) create the highest media entrainment in a population. For example, we have looked at greedily choosing nodes to be influenced by the media based on some measure of their centrality (i.e., importance [20]). We observed an increase in χ compared to the case in which we select nodes to influence uniformly at random. From this perspective, we imagine the media as advertisers, and we want to find which strategy is the most effective for advertising. Another variation is to include a competing ideology and develop algorithms that help reduce the influence of the competing ideology. We expect challenges to include showing statistically significant results from numerical computations and also analytically proving that our algorithm results in higher values of χ.

Finally, there are several modifications to the current model that we plan to explore. One possibility is extending the opinion space to dimensions $d > 1$; another is to allow the media to have different ideologies. For the latter purpose, we can draw the distribution of ideologies from empirical data. For example, we could use Figure 10, a hand-curated chart of media biases against political ideologies. Finally, we plan to make c a function of $\mathbf{x} \in \mathbb{R}^d$, as we expect people with views further from $\mathbf{0}$, who are more extreme, to be less likely than moderates to be influenced by opinions further than their own.

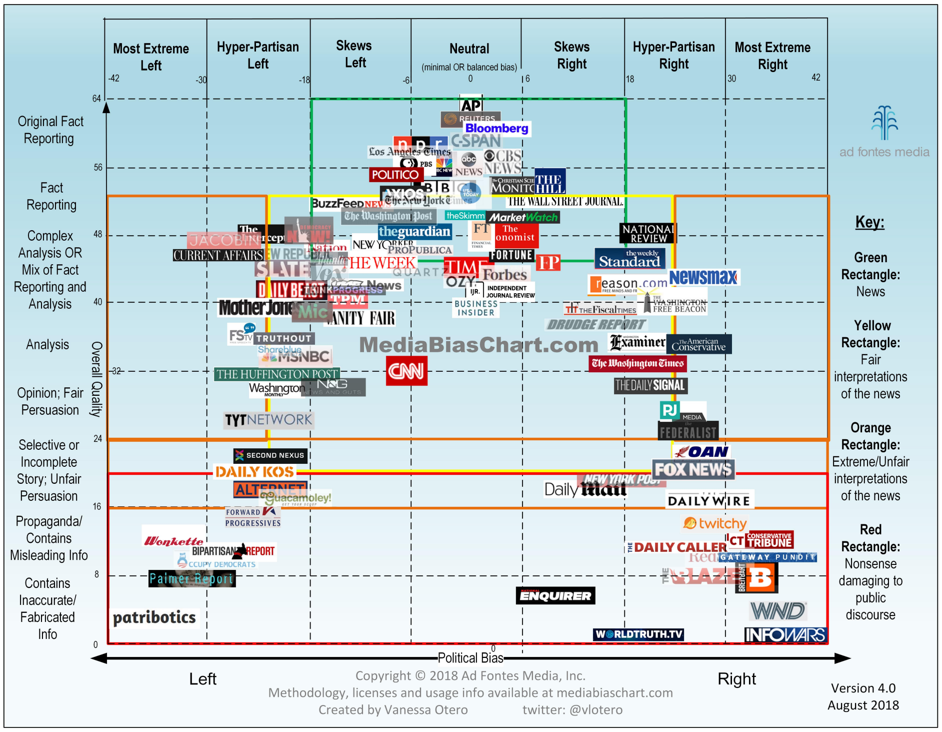

Figure 10: A chart of media bias against political ideology created from sentiment analysis of empirical data. [Taken from [18]].

References

[1] A. Bessi and E. Ferrara, *Social bots distort the 2016 U.S. Presidential election online discussion*, First Monday, 21 (2016).

[2] H. Z. Brooks and M. A. Porter, *A Model for the Influence of Media on the Ideology of Content in Online Social Networks*, arXiv e-prints, (2019), p. arXiv:1904.09238.

[3] J. A. Carrillo, Y.-P. Choi, and M. Hauray, *The derivation of swarming models: mean-field limit and Wasserstein distances*, in Collective dynamics from bacteria to crowds, vol. 553 of CISM Courses and Lect., Springer, Vienna, 2014, pp. 1–46.

[4] J. A. Carrillo, M. Di Francesco, A. Esposito, S. Fagioli, and M. Schmidtchen, *Measure solutions to a system of continuity equations driven by Newtonian nonlocal interactions*, arXiv e-prints, (2018), p. arXiv:1810.10236.

[5] J. D. Crawford, *Amplitude expansions for instabilities in populations of globally-coupled oscillators*, J. Stat. Phys., 74 (1994), pp. 1047–1084.

[6] ———, *Scaling and singularities in the entrainment of globally coupled oscillators*, Phys. Rev. Lett., 74 (1995), pp. 4341–4344.

[7] J. D. Crawford and K. Davies, *Synchronization of globally coupled phase oscillators: singularities and scaling for general couplings*, Physica D, 125 (1999), pp. 1 – 46.

[8] G. Deffuant, D. Neau, F. Amblard, and G. Weisbuch, *Mixing beliefs among interacting agents*, Adv. Complex Syst., 03 (2000), pp. 87–98.

[9] M. H. Degroot, *Reaching a consensus*, J. Amer. Statist. Assoc., 69 (1974), pp. 118–121.

[10] M. Dimock, J. Kiley, S. Keeter, and C. Doherty, *Political polarization in the american public.* `https://www.people-press.org/2014/06/12/political-polarization-in-the-american-public/`, Jun 2014. Pew Research Center.

[11] R. I. M. Dunbar, *Coevolution of neocortical size, group size and language in humans*, Behav. Brain Sci., 16 (1993), pp. 681–694.

[12] M. D. Francesco and S. Fagioli, *Measure solutions for non-local interaction PDEs with two species*, Nonlinearity, 26 (2013), pp. 2777–2808.

[13] J. R. P. French Jr., *A formal theory of social power*, Psychol. Rev., 63 (1956), pp. 181–194.

[14] J. Guckenheimer and P. Holmes, *Nonlinear oscillations, dynamical systems, and bifurcations of vector fields*, vol. 42 of Applied Mathematical Sciences, Springer-Verlag, New York, 1990. Revised and corrected reprint of the 1983 original.

[15] R. Hegselmann and U. Krause, *Opinion dynamics and bounded confidence, models, analysis and simulation*, J. Artif. Soc. Soc. Simul., 5 (2002).

[16] P.-E. Jabin and Z. Wang, *Mean Field Limit for Stochastic Particle Systems*, Springer International Publishing, Cham, Switzerland, 2017, pp. 379–402.

[17] K. E. Matsa and E. Shearer, *News use across social media platforms 2018.* `https://www.journalism.org/2018/09/10/news-use-across-social-media-platforms-2018/`, Sep 2018. Pew Research Center.

[18] A. F. Media, *Media bias chart.* `adfontesmedia.com`, Aug 2018.

[19] X. F. Meng, R. A. Van Gorder, and M. A. Porter, *Opinion formation and distribution in a bounded-confidence model on various networks*, Phys. Rev. E, 97 (2018), p. 022312.

[20] M. Newman, *Networks: An Introduction*, Oxford University Press, Inc., New York, NY, USA, 2010.

[21] S. H. Strogatz, *From Kuramoto to Crawford: Exploring the onset of synchronization in populations of coupled oscillators*, Physica D, 143 (2000), pp. 1–20.

[22] A. L. Traud, P. J. Mucha, and M. A. Porter, *Social structure of Facebook networks*, Physica A, 391 (2012), pp. 4165–4180.

[23] U. von Luxburg, *A tutorial on spectral clustering*, Stat. Comput., 17 (2007), pp. 395–416.