

Multipolar opinion evolution in biased networks

Luka Baković, David Ohlin, Giacomo Como and Emma Tegling

Abstract—Motivated by empirical research on bias and opinion formation, we formulate a multidimensional nonlinear opinion-dynamical model where agents have individual biases, which are fixed, as well as opinions, which evolve. The dimensions represent competing options, of which each agent has a relative opinion, and are coupled through normalization of the opinion vector. This can capture, for example, an individual’s relative trust in different media. In special cases including where biases are uniform across agents our model achieves consensus, but in general, behaviors are richer and capture multipolar opinion distributions. We examine general fixed points of the system, as well as special cases such as zero biases toward certain options or partitioned decision sets. Lastly, we demonstrate that our model exhibits polarization when biases are spatially correlated across the network, while, as empirical research suggests, a mixed community can mediate biases.

I. INTRODUCTION

Social media is becoming increasingly relevant as a news source, as more and more people use it either as a way of aggregating content from different sources directly or gathering information by reading articles shared by their friends. More than half of Americans get news from social media, and for millennials [1] shows that social media is even the primary news source. With the range of news sources available, a key question is how users decide whom to trust. While *accuracy* of the news is an important contributing factor to media trust [2], it is also true that perceived accuracy may be subjective. For example, left, center and right leaning news outlets all have their audiences, which would imply three differing viewpoints being accurate at the same time. Indeed, *bias* plays a key role and has long been the subject of social science research. As shown in [3], individuals have a tendency to gather evidence which is partial to their preconceptions. In other words, we perceive something as accurate if it is in line with our biases and preconceptions, see [4] for an elaboration. In this paper, we propose a novel opinion-dynamical model that can capture how preconceptions or biases impact people’s evolving trust in the news sources that their social networks expose them to.

Recent empirical research on social networks shows that they do in fact exhibit phenomena such as polarization [5] and ideological segregation [6], meaning that biases seem to play a role in how information is distributed and opinions

The authors are with the Department of Automatic Control, Lund University, Lund, Sweden. They are members of the ELLIIT Strategic Research Area at Lund University. Email: {luka.bakovic, david.ohlin, giacomo.como, emma.tegling}@control.lth.se.

Giacomo Como is also with the Department of Mathematical Sciences, Politecnico di Torino, Torino, Italy. giacomo.como@polito.it

This work was partially supported by ELLIIT and the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation.

are formed. It is clear, however, that the mapping between bias and opinions is not an identity map. Indeed, as shown in research on contact theory, outlined in [7], intergroup contact can act as a mediator for biases. Sufficiently diverse networks can imply bias mediation. This is found in [8] to be the case – greater diversity can lead to greater perceptions of similarity. This phenomenon is also captured by our model.

With this empirical context in mind, the contribution of this paper is to propose a novel multidimensional nonlinear opinion-dynamical model where agents have an underlying bias, as well as opinions, which evolve as the information shared by their neighbors in a social network is filtered through the bias. More specifically, agents in our model update their opinion by combining it with an individually weighted (i.e. biased) average of their neighbors’ opinions in each dimension. The vector components, that is, the opinions on different issues, are then coupled through a normalization that ensures the opinion vector sums to 1. The opinion vector may, for example, represent an individual’s relative trust in, or attention to, various news sources, where an increase in attention to one source will lead to a decrease in another. The nonlinearity allows the model expressivity in the sense that it exhibits different steady states ranging from polarization to consensus, depending on the configuration and initial conditions. For example, introducing spatial correlation between biases leads to a polarization of opinions. In other words, communities in the graph will become echo chambers if given similar biases, while diversity can mediate the biases.

A. Related work

Some of the most commonly studied models in opinion dynamics are linear ones which use averaging as principle mechanism for social interaction. Such are the DeGroot model [9] and the Friedkin-Johnsen model [10], [11]. A detailed analysis of both can be found in [12], [13]. Both models predict convergence to a convex combination of the starting opinions, and are therefore unsuitable for modeling of radicalization, formation of echo chambers or segregation, which real networks do exhibit as shown in [5] and [6]. Compared to them, our model is more expressive as it only reaches consensus in some specific cases and allows opinions to escape the convex hull of initial opinions. Our model allows like-minded individuals to push each other further toward extreme opinions. It is worth noting, however, that for a certain choice of parameters (uniform bias) our model reduces to the DeGroot model.

Conceptually, our model also relates to the bounded confidence model introduced in [14]. There, agents only interact with agents whose opinions are close enough to theirs, mean-

ing that the network itself evolves with the opinions. Whilst successfully capturing segregation of opinions, this model still converges to convex combinations of starting opinions. Multidimensional versions of [14] also exist, for example in [15]. A related multidimensional model is proposed by the simulation study [16], where the interaction between agents is determined by their agreement across the opinion vector, also accounting for negative influence. This is in contrast to our model, where agents with different opinions can still interact if their biases permit it. A more intricate coupling mechanism motivated by epidemical modeling is proposed by [17]. There, the vaccination network can be compared to our bias and the epidemic to the opinion. Both networks then evolve dynamically, whereas we fix the bias and consider opinion dynamics only.

In certain aspects, our model also relates to [18], wherein a linear operator is applied to the "neighborhood" term of the averaging dynamics, thus coupling the opinion vector. The nonlinearity of our model, however, captures other behaviors. The nonlinear opinion dynamics model presented in [19] has recently received a lot of attention. Our models exhibit similarities such as the ability for multi-option opinion formation and parameters mimicking preference or bias of agents toward certain issues. Both models include some simpler models as special cases. Our models differ in domain choice and type of nonlinearity, as well as the underlying empirical motivation. As opinions evolve, agents in our model can stop adapting to certain neighbors. This is not unlike the self-appraisal mechanism of [20].

B. Contributions and paper outline

This paper proposes a novel multidimensional nonlinear opinion-dynamical model, motivated by the body of empirical work on bias and preconceptions. It provides an initial analysis of its behaviors that depend on the biases, the opinions as well as the social network structure. We provide a characterization of fixed points of the system and give sufficient conditions for the special case where agents reach a consensus. We also present a Lyapunov argument for the case where one set of options is strictly preferred by all agents. We illustrate the dynamics through numerical examples, which in particular demonstrate emergence of polarization in cases where biases are spatially correlated.

The model is introduced in Section II. Section III contains analytical results on the system's fixed points, which are then illustrated through numerical examples in Section IV. We conclude with directions for future work in Section V.

II. THE MODEL

A. Notation and Preliminaries

Bold fonts indicate vectors, capital letters indicate matrices. Let $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ be a finite simple graph, where \mathcal{V} is the set of nodes and \mathcal{E} undirected edges. Simple graphs are those with no multi-edges or self-loops. A simple graph is connected if there exists a path between every pair of nodes. Let \mathcal{N}_i denote the neighborhood of agent i . Agents will be assigned corresponding opinion vectors $\mathbf{x}^i \in \mathbb{R}^k$ of which

the j -th component is denoted x_j^i . We let $\|\cdot\|_1$ denote the canonical 1-norm $\|\mathbf{x}\|_1 = \sum_{i=1}^k |x_i|$. The identity matrix of size $k \times k$ is denoted \mathbb{I}_k with subscript omitted when the size is clear from context. All inequalities apply element-wise. Vectors of the canonical orthonormal basis in \mathbb{R}^k are represented by \mathbf{e}^i where $i \in \{1, \dots, k\}$. When convenient we use the shortened notation \mathbf{x}^+ , dropping the time argument, to indicate one step of the evolution of some dynamics relative to a previous state \mathbf{x} .

B. Model Setup

Let \mathcal{G} be a simple, connected graph with n agents. Each agent i in the graph has a state $\mathbf{x}^i \in \mathbb{R}_{\geq 0}^k$ located on the $(k-1)$ -simplex, defined as follows,

Definition 1 (k-simplex):

$$\Delta^k = \{\mathbf{x} \in \mathbb{R}^{k+1} : \|\mathbf{x}\|_1 = 1, \mathbf{x} \geq \mathbf{0}\}.$$

The state \mathbf{x}^i represents the relative opinion of agent i of the k alternatives. Let $\mathcal{K} = \{1, \dots, k\}$ denote the set of alternatives corresponding to the indices of \mathbf{x}^i . Assign to each agent $i \in \mathcal{V}$ a diagonal, nonnegative matrix $R^i \in \mathbb{R}_{\geq 0}^{k \times k}$. The update rule for every $t \in \mathbb{N}$ is then given by

$$\mathbf{x}^i(t+1) = \frac{\mathbf{x}^i(t) + \sum_{j \in \mathcal{N}_i} R^i \mathbf{x}^j(t)}{\|\mathbf{x}^i(t) + \sum_{j \in \mathcal{N}_i} R^i \mathbf{x}^j(t)\|_1}, \quad (1)$$

with $\mathbf{x}^i(0) = \mathbf{x}_0^i \in \Delta^{k-1}$. The update equation consists of three main components, which we briefly explain here. The first component in the numerator is the opinion of the agent from the last time step. The second term of the numerator captures the filtering of the neighbors' opinions through agent i 's bias. The index corresponding to the largest bias is the most sensitive – agents have the easiest time moving their opinion in that direction. Note that each agent has an individual bias, in contrast to [18] where one bias between topics is defined for the whole graph.

Finally, the denominator brings the opinion vector back onto Δ^{k-1} , in other words, the opinion vector is normalized. As a result, an increase in one index necessarily causes a decrease in some other index. This is in contrast to standard multidimensional averaging and bounded confidence, where indices evolve independently. This models situations such as the choice between competing products, or probabilistic ones such as the split of an agent's attention between different news sources. The normalization is a mechanism for coupling, and permits more nuanced behaviors.

Remark 1: The bias matrix R^i is uniquely determined by a vector $\mathbf{r}^i \in \mathbb{R}^k$ such that $R^i = \text{diag}(\mathbf{r}^i)$. We will be referring to the vector in places where it aids clarity.

Remark 2: Agent $i \in \mathcal{V}$ can increase its opinion in index $\ell \in \mathcal{K}$ only if the corresponding diagonal entry of its bias matrix R^i is positive in that index and there exists an agent in its neighborhood with a positive opinion in that index.

Example 1: Let $k = 2$. This could be interpreted as a scenario of choosing between two different options. We examine an agent who is currently undecided but prejudiced toward the first option, meaning that

$$\mathbf{x}^1(0) = \begin{bmatrix} 0.5 \\ 0.5 \end{bmatrix} \quad R^1 = \begin{bmatrix} 0.6 & 0 \\ 0 & 0.4 \end{bmatrix}.$$

We introduce two neighbors with initial opinions $\mathbf{x}^2(0) = [0.8 \ 0.2]^\top$ and $\mathbf{x}^3(0) = [0.2 \ 0.8]^\top$ and calculate the sum in the numerator of the update term for \mathbf{x}^1 .

$$R^1 \mathbf{x}^2(0) + R^1 \mathbf{x}^3(0) = \begin{bmatrix} 0.48 \\ 0.08 \end{bmatrix} + \begin{bmatrix} 0.12 \\ 0.32 \end{bmatrix}$$

After normalization, the updated opinion becomes

$$\mathbf{x}^1(1) = \begin{bmatrix} 0.55 \\ 0.45 \end{bmatrix}.$$

The neighbors were equally sure about their options, but $\mathbf{x}^2(0)$ was more in line with the agents's bias; thus it produced a comparatively greater push toward its option. Further iterations and the steady state of the model depend both on R^2 , R^3 and the structure of the interaction graph, as can be seen from the examples in Section IV.

Next, let us give a simplistic real-world example describing the role of the bias vector \mathbf{r} in the opinion evolution.

Example 2: Consider a case where agents represent individuals whose incomes can be roughly divided in a lower, middle and upper bracket. This might be encoded in their bias vectors with $[1 \ 0 \ 0]$ corresponding to a lower bracket, $[0.5 \ 0.5 \ 0]$ a lower-middle bracket and so on. Agents then form an opinion on which bracket deserves a tax cut in the budget. Their opinion may then align with their bias, but only if presented with that option by their neighbors, that is, if the neighborhood is fairly homogeneous. An upper middle class earner $\mathbf{r}^i = [0 \ 0.6 \ 0.4]$ surrounded by upper class neighbors $\mathbf{r}^j = [0 \ 0 \ 1]$, on the other hand, has a high chance of advocating for the upper class as theirs is the locally dominant opinion. Trust in media with different political leanings may follow a similar logic.

The following proposition shows that, for certain parameter choices, our model reduces to a simple, well-known averaging model.

Proposition 1: For any $k \in \mathbb{N}$, $n \in \mathbb{N}$ and \mathcal{G} satisfying the definition, setting $R^i = \mathbb{I}_k$ ($\forall i \in \mathcal{V}$) reduces the model to the DeGroot model [9].

Proof: Inserting in $R^i = \mathbb{I}_k$ in (1) gives

$$\mathbf{x}^i(t+1) = \frac{\mathbf{x}^i(t) + \sum_{j \in \mathcal{N}_i} \mathbf{x}^j(t)}{\|\mathbf{x}^i(t) + \sum_{j \in \mathcal{N}_i} \mathbf{x}^j(t)\|_1}.$$

Since every \mathbf{x}^i is nonnegative and unit in the 1-norm, the denominator simplifies, leaving the averaging formula:

$$\mathbf{x}^i(t+1) = \frac{\mathbf{x}^i(t) + \sum_{j \in \mathcal{N}_i} \mathbf{x}^j(t)}{1 + |\mathcal{N}_i|}.$$

Remark 3: In Proposition 1, setting biases to $c \mathbb{I}_k$ for any $c \in \mathbb{R}_+$ gives the same result, the only difference being the convergence rate. ■

III. ANALYSIS

In this section, we present certain analytical results pertaining to the fixed points of the system. While these are difficult to characterize completely due to the possible diversity of outcomes depending on the network topology, bias vectors and initial opinions, we give conditions for a number of relevant special cases. The objective is to characterize

the mathematical properties of the model itself, to inform its applicability in practical contexts. The following is a technical result on trivial fixed points.

Proposition 2: Configurations where all agents share the same state with a 1 in one entry and a 0 in all other entries

$$\mathbf{x}^1, \dots, \mathbf{x}^n = [0, \dots, 0, 1, 0, \dots, 0]^\top$$

are fixed points of (1).

Proof: Suppose $\mathbf{x}^i = \mathbf{e}^\ell$ for all $i \in \mathcal{V}$ and some $\ell \in \mathcal{K}$.

The result follows from insertion in (1). ■

Example 3: These fixed points can easily be shown to be attractive for certain trivial configurations of the model. Let $k \in \mathbb{N}$ and $n \in \mathbb{N}$. Then, if there exists an index ℓ such that $R^i = \text{diag}(\mathbf{e}^\ell)$ for all agents $i \in \mathcal{V}$, it holds that $\mathbf{x}^i(t) \rightarrow \mathbf{e}^\ell$ for all $i \in \mathcal{V}$ and for almost all initial conditions.

A. General fixed points and bias with zero elements

We next give a general result characterizing the fixed points of (1). In cases where zero bias toward some options is allowed, agents may stop communicating with neighbors that have certain opinions. This in turn allows for fixed points where the opinions of some agents are not uniquely determined by neighbors' opinions.

Proposition 3: At the fixed points of the dynamics (1), the state \mathbf{x}^i corresponding to an agent $i \in \mathcal{V}$ satisfies exactly one of the following two conditions:

- (i) $\|\sum_{j \in \mathcal{N}_i} R^i \mathbf{x}^j\|_1 = 0$
- (ii) $\mathbf{x}^i = \frac{1}{\|\sum_{j \in \mathcal{N}_i} R^i \mathbf{x}^j\|_1} \sum_{j \in \mathcal{N}_i} R^i \mathbf{x}^j$

Proof: For all fixed points of (1) it must hold that $(\mathbf{x}^i)^+ = \mathbf{x}^i$, which is equivalent to

$$\frac{\mathbf{x}^i(1 - \|\mathbf{x}^i + \sum_{j \in \mathcal{N}_i} R^i \mathbf{x}^j\|_1) + \sum_{j \in \mathcal{N}_i} R^i \mathbf{x}^j}{\|\mathbf{x}^i + \sum_{j \in \mathcal{N}_i} R^i \mathbf{x}^j\|_1} = 0$$

$$\iff \mathbf{x}^i \|\sum_{j \in \mathcal{N}_i} R^i \mathbf{x}^j\|_1 = \sum_{j \in \mathcal{N}_i} R^i \mathbf{x}^j,$$

where division by $\|\sum_{j \in \mathcal{N}_i} R^i \mathbf{x}^j\|_1$ yields the expression (i) for \mathbf{x}^i at the fixed point unless $\|\sum_{j \in \mathcal{N}_i} R^i \mathbf{x}^j\|_1 = 0$. ■

Corollary 4: Let the bias \mathbf{r}^i of some agent i have a zero element. Then, at the fixed points of (1), either $\mathbf{x}_\ell^i = 0$ or (i) in Proposition 3 holds for this agent.

Corollary 4 is a direct consequence of Proposition 3 in the case where agents have no susceptibility to some option ℓ . It means that in any fixed configuration the agents either reject the option they are not susceptible to ($\mathbf{x}_{\ell_i} = 0$), or they are no longer influenced by their neighbors.

Corollary 5: Let \mathbf{r}^i have strictly positive entries for all agents i . At the fixed points of (1), if any agent is on the boundary of Δ^{k-1} (that is, $\mathbf{x}_\ell^i = 0$ for at least one index $\ell \in \mathcal{K}$) then all other agents must be at the same face.

Proof: Condition (i) in Proposition 3 can never hold under this assumption on the biases \mathbf{r}^i . Evaluating the expression in point (ii) of Proposition 3 we see that \mathbf{x}_ℓ^i can only be zero at a fixed point if all neighbors $j \in \mathcal{N}_i$ have $\mathbf{x}_\ell^j = 0$. Repeating the argument for these neighbors yields that $\mathbf{x}_\ell^i = 0$ must be true for all agents. ■

B. Rejection of non-dominant alternatives

In order to present the next result, we first introduce a partition of the set of alternatives \mathcal{K} into dominant and recessive opinions.

Definition 2: Let \mathcal{K} be partitioned into two subsets; the dominant set \mathcal{D} and the recessive set \mathcal{L} . The recessive set is defined as the largest set $\mathcal{L} \subsetneq \mathcal{K}$ such that for all $\ell \in \mathcal{L}$ and $d \in \mathcal{K} \setminus \mathcal{L}$ the inequality $r_\ell^i < r_d^i$ holds for all agents i . The dominant set is defined as the remainder, $\mathcal{D} = \mathcal{K} \setminus \mathcal{L}$.

Remark 4: Note that the largest set \mathcal{L} is uniquely defined (though it may be empty) and all smaller sets fulfilling the criteria are subsets of \mathcal{L} .

We can now show that opinions in the set \mathcal{L} are suppressed by the dynamics.

Proposition 6: At the fixed points of (1), for almost all initial values, all agents i have opinion $\mathbf{x}_\ell^i = 0$ for $\ell \in \mathcal{L}$.

Proof: We define the Lyapunov candidate

$$V(\mathbf{x}) = \max_i \sum_{\ell \in \mathcal{L}} x_\ell^i \quad (2)$$

which satisfies $V(\mathbf{x}) = 0$ only when, for every agent i , $x_\ell^i = 0$ for all $\ell \in \mathcal{L}$. We wish to show that

$$V(\mathbf{x}^+) - V(\mathbf{x}) < 0 \quad (3)$$

on the domain

$$\{\mathbf{x} : \mathbf{x}^i \in \Delta^{k-1}, V(\mathbf{x}) \neq 1\} \quad (4)$$

which is invariant under the dynamics (1). Let $\bar{x}_\mathcal{L} := \max_i (\sum_{\ell \in \mathcal{L}} x_\ell^i)$. Inserting (1) into (3) gives

$$\begin{aligned} V(\mathbf{x}^+) - V(\mathbf{x}) &= \max_i \sum_{\ell \in \mathcal{L}} \left(\frac{x_\ell^i + \sum_{j \in \mathcal{N}_i} r_\ell^i x_\ell^j}{1 + \|\sum_{j \in \mathcal{N}_i} R^i \mathbf{x}^j\|_1} \right) - \bar{x}_\mathcal{L} \\ &\leq \max_i \frac{\bar{x}_\mathcal{L} + \sum_{\ell \in \mathcal{L}} \sum_{j \in \mathcal{N}_i} r_\ell^i x_\ell^j}{1 + \|\sum_{j \in \mathcal{N}_i} R^i \mathbf{x}^j\|_1} - \bar{x}_\mathcal{L} \\ &= \max_i \frac{\sum_{\ell \in \mathcal{L}} \sum_{j \in \mathcal{N}_i} r_\ell^i x_\ell^j - \bar{x}_\mathcal{L} \|\sum_{j \in \mathcal{N}_i} R^i \mathbf{x}^j\|_1}{1 + \|\sum_{j \in \mathcal{N}_i} R^i \mathbf{x}^j\|_1}. \end{aligned}$$

If we can show that the expression being maximized is strictly negative for all i then (3) must also hold. This yields the condition $\sum_{\ell \in \mathcal{L}} \sum_{j \in \mathcal{N}_i} r_\ell^i x_\ell^j < \bar{x}_\mathcal{L} \|\sum_{j \in \mathcal{N}_i} R^i \mathbf{x}^j\|_1$. A more strict condition, implying the above, is given by the element-wise expression $\sum_{\ell \in \mathcal{L}} r_\ell^i x_\ell^j < \bar{x}_\mathcal{L} \|R^i \mathbf{x}^j\|_1$. Let $\rho^i = \min_{\ell \in \mathcal{D}} r_\ell^i$. Division by ρ^i yields

$$\begin{aligned} \sum_{\ell \in \mathcal{L}} \frac{r_\ell^i}{\rho^i} x_\ell^j &< \frac{1}{\rho^i} \bar{x}_\mathcal{L} \|R^i \mathbf{x}^j\|_1 \\ \iff (1 - \bar{x}_\mathcal{L}) \sum_{\ell \in \mathcal{L}} \frac{r_\ell^i}{\rho^i} x_\ell^j &< \bar{x}_\mathcal{L} \sum_{\ell \in \mathcal{D}} \frac{r_\ell^i}{\rho^i} x_\ell^j. \end{aligned} \quad (5)$$

Note that according to Definition 2 we have $r_\ell^i/\rho^i < 1$ for $\ell \in \mathcal{L}$ and $r_\ell^i/\rho^i \geq 1$ for $\ell \in \mathcal{D}$. As a consequence of the definition of $\bar{x}_\mathcal{L}$, the factors of (5) satisfy

$$(1 - \bar{x}_\mathcal{L}) \leq \sum_{\ell \in \mathcal{D}} \frac{r_\ell^i}{\rho^i} x_\ell^j \quad \text{and} \quad \sum_{\ell \in \mathcal{L}} \frac{r_\ell^i}{\rho^i} x_\ell^j < \bar{x}_\mathcal{L},$$

for all $\mathbf{x}^i \in \Delta^{k-1}$ such that $\sum_{\ell \in \mathcal{L}} x_\ell^i \neq 1$, thus implying (5). This proves that V is strictly decreasing on the domain stated in (4), concluding the proof. \blacksquare

Proposition 6 is a versatile result, giving formal motivation for several intuitive properties of the model. The most obvious is the case of aligned biases, where $\mathbf{r}^i = \mathbf{r}$ (or when \mathbf{r} has the same ordering of alternatives). In this case the largest proper subset \mathcal{L} in Definition 2 contains all alternatives except the one with highest bias. Application of Proposition 6 to this instance yields a consensus for the alternative favored by the biases, which is the natural conclusion. Further, in cases where all agents are biased against a certain set of alternatives \mathcal{L} , these are suppressed by the dynamics and not present in the final opinion distribution. This happens despite the fact that the agents may not agree on the best option, and regardless of whether some agents initially held a favorable opinion of the alternatives in \mathcal{L} .

C. Fixed point stability analysis for $n = 2$, $k = 2$

Next, we examine in further detail the simplest non-trivial case of $n = 2$ agents and opinion dimension $k = 2$. Due to a reasonable number of parameters, this scenario lends itself to direct analysis. Now, since $x_1^1 + x_2^1 = 1$, examining one component per agent is sufficient and we let x^i denote x_1^i . Let $\mathbf{r}^1 = [\alpha^1 \beta^1]^\top$ and $\mathbf{r}^2 = [\alpha^2 \beta^2]^\top$. The dynamics (1) become:

$$x^1(t+1) = \frac{x^1(t) + \alpha^1 x^2(t)}{1 + \beta^1 + (\alpha^1 - \beta^1)x^2(t)} \quad (6a)$$

$$x^2(t+1) = \frac{x^2(t) + \alpha^2 x^1(t)}{1 + \beta^2 + (\alpha^2 - \beta^2)x^1(t)} \quad (6b)$$

of which the fixed points are given by:

$$x^{1*} = \frac{\alpha^1 x^{2*}}{\alpha^1 x^{2*} + \beta^1(1 - x^{2*})} \quad (7a)$$

$$x^{2*} = \frac{\alpha^2 x^{1*}}{\alpha^2 x^{1*} + \beta^2(1 - x^{1*})}. \quad (7b)$$

Substituting (7) into (7a) gives a quadratic equation for the fixed points of agent 1:

$$(\alpha^1 \alpha^2 - \beta^1 \beta^2) ((x^{1*})^2 - x^{1*}) = 0. \quad (8)$$

To determine local asymptotic stability of the fixed points, we will make use of the following lemma.

Lemma 7: A nonnegative matrix $\begin{bmatrix} a & b \\ c & d \end{bmatrix}$ is Schur stable if and only if $a + d < 2$ and $a + d + bc < 1 + ed$.

Proof: By the Perron-Frobenius theorem, the dominant eigenvalue is real and positive. Explicit computation gives $\lambda_{\max} = \frac{1}{2} (a + d + \sqrt{(a + d)^2 - 4ad + 4bc})$. Finally, $\lambda_{\max} < 1$ is equivalent to the statement in the lemma. \blacksquare

We can now characterize the fixed points for this low-dimensional case:

Proposition 8: The fixed points of (6a)–(6b) depend on the bias vectors $\mathbf{r}^1 = [\alpha^1 \beta^1]^\top$ and $\mathbf{r}^2 = [\alpha^2 \beta^2]^\top$ as follows:

- (i) if $\alpha^1 \alpha^2 < \beta^1 \beta^2$ then (0,0) is a locally asymptotically stable fixed point while (1,1) is an unstable fixed point,
- (ii) if $\alpha^1 \alpha^2 > \beta^1 \beta^2$ then (1,1) is a locally asymptotically stable fixed point while (0,0) is an unstable fixed point,




Fig. 1: The $n = k = 2$ case, despite its simplicity, showcases different fixed points depending on the biases and starting opinions of the agents. Biasing agents completely toward opposing options leads to a race to the boundary of the domain, as seen in (a). The symmetric but non-maximal biases in (b) give fixed points in the interior of the simplex. Finally, (c) showcases that agents with neutral biases can be influenced toward either option.

(iii) if $\alpha^1\alpha^2 = \beta^1\beta^2$ then there exists a continuum of fixed points within Δ^1 .

Proof: The fixed points in cases (i) and (ii) are easily obtained from (8). To determine stability, examine the Jacobian of the system at $(0,0)$:

$$\mathbf{J}(0,0) = \begin{bmatrix} \frac{1}{1+\beta^1} & \frac{\alpha^1}{1+\beta^1} \\ \frac{\alpha^2}{1+\beta^2} & \frac{1}{1+\beta^2} \end{bmatrix}.$$

Applying Lemma 7 gives that the fixed point $(0,0)$ is locally asymptotically stable if and only if $\alpha^1\alpha^2 < \beta^1\beta^2$. By symmetry, $(1,1)$ is stable if and only if $\alpha^1\alpha^2 > \beta^1\beta^2$. In case (iii) (8) reduces to $0 = 0$ implying no additional criteria for the fixed points other than (7a) and (7). ■

In short, agents converge toward their combined preference unless they, as a pair, happen to be equally predisposed toward both options. Numerical evidence suggests that larger networks exhibit similar phenomena, where the choice of biases affects the cardinality of the set of fixed points.

IV. NUMERICAL EXAMPLES

This section presents two numerical examples; the first to provide intuition of the model's behaviors through the simple two-agent two-dimensional case, the second to illustrate the interesting phenomenon of polarization versus bias mediation in a larger network.

A. Two Agents

In order to illustrate the behaviors described in Section III-C for the smallest non-trivial case, we provide three examples with qualitatively different behaviors.

1) $\alpha^1\alpha^2 = \beta^1\beta^2$

In line with Proposition 8, these combinations of biases admit a continuum of fixed points. First, in Fig. 1a, an instance with $\mathbf{r}^1 = [1 \ 0]^\top$ and $\mathbf{r}^2 = [0 \ 1]^\top$ is shown. The agents' preconceptions are completely incompatible, and their opinions move toward their biases and away from each other. This also serves as an illustration of Corollary 4, where zeros in bias vectors result in fixed points on the boundary or agents ceasing to influence each other.

Fig. 1b showcases similar dynamics, the difference being that a choice of non-maximal yet symmetrical biases $\mathbf{r}^1 = [0.7 \ 0.3]^\top$ and $\mathbf{r}^2 = [0.3 \ 0.7]^\top$ results in a fixed point in the interior of the simplex. The two agents mediate, each

moving their opinion vector toward their bias vector but stopping before reaching the boundary.

2) $\alpha^1\alpha^2 \neq \beta^1\beta^2$

Moving the bias of agent 2 closer to that of agent 1 further changes the behavior of the model. Fig. 1c shows a simulation with $\mathbf{r}^1 = [0.7 \ 0.3]^\top$ and $\mathbf{r}^2 = [0.45 \ 0.55]^\top$. Having a relatively neutral bias, agent 2 is easily influenced by its neighbor. Applying Proposition 8, we see that, since $\alpha^1\alpha^2 = 0.7 \cdot 0.45 < 0.3 \cdot 0.55 = \beta^1\beta^2$, $(1,1)$ is the locally asymptotically stable fixed point. This agrees with the simulation, as both agents converge to option 1.

B. Spatial Correlation and Polarization

For a larger-scale example, consider a connected Watts-Strogatz network [21] as the underlying graph. A characteristic of such graphs is the formation of communities – subsets of agents where intragroup connections are stronger than those with the rest of the graph. Supposing they are given similar biases, such communities can turn into echo chambers under our model, effectively reinforcing each others' biases and ignoring the majority opinion.

To run a numerical experiment, a Watts-Strogatz graph of $n = 500$ agents was generated using NetworkX and a community detection algorithm was applied. This revealed three communities of size 226, 222, and 52 within the particular graph. Initial opinions of dimension $k = 3$ were uniformly drawn from the simplex, as can be seen in Fig. 2a. Two scenarios were then analyzed.

1) *Spatially correlated biases*

All agents in the two larger communities are given a bias of $\mathbf{r}_A = [0.8 \ 0.09 \ 0.11]^\top$. This is a way of forming a "majority configuration" in the graph. Agents in the smallest community are given a bias of $\mathbf{r}_B = [0.11 \ 0.09 \ 0.8]^\top$.

As can be seen from the biases, communities agree in having least preference for option 2, but disagree in which option is most favorable. The steady state, as can be seen in Fig. 2b, is a polarization of opinions. The majority groups go toward option 1, and the minority goes toward option 2. Note that the steady state is not a simple split consensus, as final opinions are actually spread around the corner points of the simplex.

2) *Randomly assigned biases*

Once again, a majority and a minority group are chosen, this time completely at random. The group sizes match with




Fig. 2: Networks with communities such as Watts-Strogatz networks can showcase interesting behaviors. A uniform starting distribution of opinions as in (a) can result in a polarized steady state (b) when biases are spatially correlated, that is, if certain small communities are assigned a bias that is opposite to the majority. On the other hand, assigning the same number of contrarian biases to individuals chosen uniformly at random across the network leads to agreement, as can be seen in (c). The figures display the results of a simulation done on a $n = 500$ agent network.

those in the previous scenario, group A with 448 agents and group B with 52 agents. As in the previous scenario, groups are assigned biases r_A and r_B , respectively. The steady state can be seen in Fig. 2c. Agents in the steady state agree in the sense of choosing an option, but stay spread around the corner of the simplex instead of reaching consensus. Individual biases still play a role in the final opinion, albeit a comparatively smaller one than in the previous scenario. The interpretation here is that the majority manages to integrate the minority, as they are no longer in an echo chamber but are well connected to the rest of the network. While these behaviors are not surprising in light of the model’s mechanisms, our model offers a framework for deriving and describing these known behaviors.

V. CONCLUSIONS

We have proposed an opinion-dynamical model that describes individuals’ relative opinion of multiple coupled alternatives, depending on an individual bias, and connections over a social network. Despite the apparent simplicity of the model (1), it is expressive, as showcased through our numerical examples. In particular, simulation on the Watts-Strogatz network manages to capture the polarization and integration phenomena appearing in the social science literature that motivated the model. We analytically show sensible properties in choice-making scenarios, such as the rejection of alternatives if they are unfavorable to all agents.

There exist several directions for future work. First, the model can be extended in various directions; e.g. to consider a full bias matrix, allowing further coupling of different options, or a weighted graph to allow for finer specification of relationships between agents. Another direction is stochasticity in opinion exchanges, in order to adjust for imperfect communication often occurring in real discussions. Taking the initial analysis on fixed points in this work as a starting point, we are also interested in further classification of the set of fixed points and their stability in a general case, along with concrete results on convergence of the model. Finally, the validation of the model on empirical data on trust in government and media is ongoing.

ACKNOWLEDGEMENT

We thank Mark Jeeninga, Jonas Hansson and Richard Pates for fruitful discussions related to this work.

REFERENCES

- [1] Eurobarometer, “Standard Eurobarometer 92, media use in the European Union,” 2019. [Online]. Available: <https://europa.eu/eurobarometer/surveys/detail/2255>
- [2] M. Schranz, J. Schneider, and M. Eisenegger, *Media Trust and Media Use*. Springer Fachmedien Wiesbaden, 2018, p. 73–91.
- [3] R. S. Nickerson, “Confirmation bias: A ubiquitous phenomenon in many guises,” *Rev Gen Psychol*, vol. 2, no. 2, p. 175–220, June 1998.
- [4] S. Mullainathan and A. Shleifer, “The market for news,” *Am Econ Rev*, vol. 95, no. 4, pp. 1031–1053, 2005.
- [5] A. Urman, “Context matters: political polarization on Twitter from a comparative perspective,” *Media Cult and Soc*, vol. 42, no. 6, p. 857–879, Oct. 2019.
- [6] González-Bailón. et al., “Asymmetric ideological segregation in exposure to political news on Facebook,” *Science*, vol. 381, no. 6656, p. 392–398, July 2023.
- [7] G. W. Allport, *The nature of prejudice*. Cambridge, MA: Addison-Wesley, 1954.
- [8] X. Bai, M. R. Ramos, and S. T. Fiske, “As diversity increases, people paradoxically perceive social groups as more similar,” *Proc Natl Acad Sci*, vol. 117, no. 23, p. 12741–12749, May 2020.
- [9] M. H. DeGroot, “Reaching a consensus,” *J Am Stat Assoc*, vol. 69, no. 345, pp. 118–121, 1974.
- [10] N. E. Friedkin and E. C. Johnsen, “Social influence and opinions,” *J Math Sociol*, vol. 15, no. 3-4, pp. 193–206, 1990.
- [11] —, “Social influence networks and opinion change,” *Adv Group Process*, vol. 16, no. 1, pp. 1–29, 1999.
- [12] A. V. Proskurnikov and R. Tempo, “A tutorial on modeling and analysis of dynamic social networks. Part I,” *Annu Rev Control*, vol. 43, pp. 65–79, 2017.
- [13] —, “A tutorial on modeling and analysis of dynamic social networks. Part II,” *Annu Rev Control*, vol. 45, pp. 166–190, 2018.
- [14] G. Deffuant, D. Neau, F. Amblard, and G. Weisbuch, “Mixing beliefs among interacting agents,” *Adv Complex Syst*, vol. 03, no. 01n04, pp. 87–98, Jan. 2000.
- [15] J. Lorenz, *Fostering Consensus in Multidimensional Continuous Opinion Dynamics under Bounded Confidence*. Berlin Heidelberg: Springer, 2008, pp. 321–334.
- [16] S. Schweighofer, D. Garcia, and F. Schweitzer, “An agent-based model of multi-dimensional opinion dynamics and opinion alignment,” *Chaos*, vol. 30, no. 9, 2020.
- [17] Peng et al., “A multilayer network model of the coevolution of the spread of a disease and competing opinions,” *Math. Models Methods Appl. Sci.*, vol. 31, no. 12, pp. 2455–2494, Nov. 2021.
- [18] S. E. Parsegov, A. V. Proskurnikov, R. Tempo, and N. E. Friedkin, “Novel multidimensional models of opinion dynamics in social networks,” *IEEE Trans Autom Control*, vol. 62, no. 5, pp. 2270–2285, May 2017.
- [19] A. Bizyaeva, A. Franci, and N. E. Leonard, “Nonlinear opinion dynamics with tunable sensitivity,” *IEEE Trans Autom Control*, vol. 68, no. 3, p. 1415–1430, Mar. 2023.
- [20] P. Jia, A. MirTabatabaei, N. E. Friedkin, and F. Bullo, “Opinion dynamics and the evolution of social power in influence networks,” *SIAM REV*, vol. 57, no. 3, p. 367–397, Jan. 2015.
- [21] D. J. Watts and S. H. Strogatz, “Collective dynamics of ‘small-world’ networks,” *Nature*, vol. 393, no. 6684, pp. 440–442, June 1998.