

Bounded-confidence model of opinion dynamics with heterogeneous node-activity levelsGrace J. Li *Department of Mathematics, University of California, Los Angeles, California 90095, USA*Mason A. Porter **Department of Mathematics, University of California, Los Angeles, California 90095, USA
and Santa Fe Institute, Santa Fe, New Mexico 87501, USA*(Received 19 June 2022; accepted 22 March 2023; published 21 June 2023;
corrected 30 June 2023 and 4 March 2024)

Agent-based models of opinion dynamics allow one to examine the spread of opinions between entities and to study phenomena such as consensus, polarization, and fragmentation. By studying models of opinion dynamics on social networks, one can explore the effects of network structure on these phenomena. In social networks, some individuals share their ideas and opinions more frequently than others. These disparities can arise from heterogeneous sociabilities, heterogeneous activity levels, different prevalences to share opinions when engaging in a social-media platform, or something else. To examine the impact of such heterogeneities on opinion dynamics, we generalize the Deffuant-Weisbuch (DW) bounded-confidence model (BCM) of opinion dynamics by incorporating node weights. The node weights allow us to model agents with different probabilities of interacting. Using numerical simulations, we systematically investigate (using a variety of network structures and node-weight distributions) the effects of node weights, which we assign uniformly at random to the nodes. We demonstrate that introducing heterogeneous node weights results in longer convergence times and more opinion fragmentation than in a baseline DW model. The node weights in our BCM allow one to consider a variety of sociological scenarios in which agents have heterogeneous probabilities of interacting with other agents.

DOI: [10.1103/PhysRevResearch.5.023179](https://doi.org/10.1103/PhysRevResearch.5.023179)**I. INTRODUCTION**

Humans are connected in numerous ways, and our many types of interactions with each other influence what we believe and how we act. To model how opinions spread between people or other agents, researchers across many disciplines have developed a variety of models of opinion dynamics [1–7]. However, in part because of the difficulty of gathering empirical data on opinions, much of the research on opinion dynamics has focused on theory and model development, with little empirical validation [1,6–8]. Some researchers have examined how human opinions change in controlled experimental settings with questionnaires [9–11], and others have examined empirical opinion dynamics using data from social-media platforms [12–14]. One of the many difficulties in empirically validating models of opinion dynamics is the potential sensitivity of model outcomes to measurement errors of real-life opinion values [15]. See Mäs [16] for a discussion of some of the challenges of validating models in the social sciences. Even with the difficulty of validating models of opinion dynamics, it is valuable to formulate and study such

models. Developing mechanistic models forces researchers to clearly define assumptions, variables, and the relationships between variables; such models provide frameworks to explore and generate testable hypotheses about complex social phenomena [8,17].

In an agent-based model (ABM) of opinion dynamics, each agent is endowed with an opinion and an underlying network structure governs which agents can interact with each other. We assume that all interactions are dyadic (i.e., between exactly two agents), and we suppose that agent opinions take continuous values in a closed interval on the real line [18]. This interval represents a continuous spectrum of views about something, such as an ideology or the strength of support for a political candidate. At each time step of an ABM of opinion dynamics, one selects which agents interact and then an update rule determines if and how their opinions change. Bounded-confidence models (BCMs) are a popular type of ABM with continuous-valued opinions [4]. In a BCM, interacting agents influence each other only when their opinions are sufficiently similar. This mechanism is reminiscent of the psychological idea of selective exposure, which asserts that people tend to seek information or conversations that support their existing views and avoid those that challenge their views [19]. Under this assumption, an agent's views are influenced directly only by agents with sufficiently similar views. For example, social-media platforms can have polarizing posts, but individuals can choose whether or not to engage with such content. They are not persuaded by everything in their social-media feeds.

*mason@math.ucla.edu

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

The two most popular BCMs are the Hegselmann-Krause (HK) model [20] and the Deffuant-Weisbuch (DW) model [21]. At each discrete time, the HK model has synchronous updates of agent opinions. By contrast, the DW model has asynchronous opinion updates, with a single pair of agents (i.e., a dyad) interacting and potentially updating their opinions at each time. An asynchronous mechanism is consistent with empirical studies, which suggest that individuals in social networks have different activity patterns and frequencies [22]. In the present paper, we generalize the DW model to incorporate heterogeneous node-activity levels. Although the DW model has been generalized in many ways [5], few studies have modified the procedure to select which agents interact in a time step. The ones that have modified this procedure (see, e.g., Refs. [22–25]) have focused on specific scenarios, rather than on investigating the effects of introducing heterogeneities into agent-selection probabilities.

Before we describe previous extensions of the DW model that incorporate heterogeneities in agent selection, we first discuss other generalizations of the model. The DW model was first studied on complete graphs and a square-lattice graph, with the nodes of a graph representing agents [21]. To explore the effects of network structure on DW dynamics, many researchers subsequently examined DW models on a variety of time-independent graphs [26]. Researchers have also examined DW models on hypergraphs [27] and on networks that coevolve with agent opinions [28]. Additionally, many studies have extended the DW model to consider different initial conditions and/or model parameters. Some studies have considered initial agent opinions that arise from nonuniform distributions [27,29–31], yielding initial conditions that are different than those in the standard DW model. Other investigations have incorporated heterogeneous confidence bounds or heterogeneous opinion compromises [31–38]. Such generalizations affect the opinion updates of interacting agents.

In the standard DW model, one uniformly randomly selects pairs of agents to interact, but social interactions in real life are not uniformly random. Some studies of DW models have modified the selection procedure that determines which agents interact with each other [22–25]. When selecting agents in a way that is not uniformly at random, one can think of the agents as having different activity levels that encode their interaction frequencies. (In a given time interval, we expect these agents to have different numbers of interactions.) The idea of heterogeneous node-activity levels plays an important role in activity-driven models of temporal networks [39]. Activity-driven frameworks have also been used to model which agents can interact with each other in studies of opinion dynamics. Li *et al.* [40] developed an activity-driven model of opinion dynamics on networks with nodes with assigned activity rates (i.e., assigned activation probabilities). At each time step of their model, one removes all existing edges and then the active agents randomly form connections to other agents. All agents then evaluate the mean opinions of their neighbors to determine if and how to update their own opinions [40]. Researchers have also incorporated heterogeneous agent selection in voter models of opinion dynamics. Masuda *et al.* [41] studied a voter model with heterogeneous “flip” rates, which one can interpret as heterogeneous node weights that encode activity levels. Baronchelli *et al.* [42] studied a

voter model with heterogeneous edge weights, which one can interpret as encoding heterogeneous edge activities.

Some researchers have generalized the DW model to incorporate heterogeneous agent selection. Alizadeh and Cioffi-Revilla [22] studied a modified DW model that incorporates a repulsion mechanism (which was proposed initially by Huet *et al.* [43]) in which interacting agents with opinions that differ by more than a cognitive-dissonance threshold move farther away from each other in the space of opinions. They used two-dimensional (2D) vector-valued opinions and placed their nodes on complete graphs. To model agents with different activity levels, Alizadeh and Cioffi-Revilla [22] implemented a Poisson node-selection probability, which one can interpret as independent internal “clocks” that determine agent activation. In comparison to selecting agent pairs uniformly at random (as in the standard DW model), the Poisson node-selection probability can either lessen or promote the spread of extremist opinions, depending on which opinions are more prevalent in more-active agents.

Zhang *et al.* [23] examined a modified DW model with asymmetric updates on activity-driven networks. In their model, each node has a fixed activity potential, which one assigns uniformly at random from a distribution of activity potentials. The activity potential of an agent gives its probability of activating. At each time step, each active agent i randomly either (1) creates a message (e.g., a social-media post) or (2) boosts a message that was created by a neighboring agent j . If agent i boosts a message from agent j , then i updates its opinion using the standard DW update mechanism. Zhang *et al.* [23] simulated their model on a social network from Tencent Weibo (腾讯微博) and found that the distribution of activity potentials influences the location of the transition between opinion consensus and fragmentation. The node weights in our BCM are similar in spirit to the activity potentials of Zhang *et al.* [23]; they can encode the social activity levels of individuals, such as their frequencies of posting or commenting on social media. However, the way that we incorporate node weights in our BCM differs fundamentally from Ref. [23]. We consider a time-independent network G , and we select a single pair of neighboring agents to interact at each time step. We first randomly select one agent with a probability that is proportional to its node weight, and then we randomly select one of its neighbors with a probability that depends on that neighbor’s node weight. The two selected agents then update their opinions using the DW update mechanism.

Heterogeneities in which interactions occur in a social network arise not only because some individuals are more likely to have interactions, but also because some pairs of individuals are more likely to interact than other pairs [42]. The curation of content in social-media feeds is affected by homophily, which is the idea that individuals have a tendency to connect with others that are similar to themselves (e.g., perhaps they have similar ideas or beliefs) [44]. Social-media feeds tend to show content to users that aligns with their profiles and past activities [45]. To examine the effect of such algorithmic bias on opinion dynamics, Sîrbu *et al.* [24] studied a modified DW model that includes a homophily-promoting activation mechanism. At each time step, one agent is selected uniformly at random, and then one of its neighbors is selected

with a probability that depends on the magnitude of the opinion difference between that neighbor and the first agent. The simulations by Sírbu *et al.* of their model on complete graphs suggest that more algorithmic bias yields slower convergence times and more opinion fragmentation [24]. Pansanella *et al.* [25] examined the same algorithmic-bias model on a variety of networks [specifically, Erdős-Rényi, Barabási-Albert, and Lancichinetti-Fortunato-Radicchi (LFR) graphs], and they found similar trends as Sírbu *et al.* did on complete graphs.

From the investigations in Refs. [22–25], we know that incorporating heterogeneous agent-selection probabilities into a DW model can influence opinion dynamics. Each of these papers examined a specific implementation of heterogeneous agent selection. We are not aware of any systematic investigations of the effects of heterogeneous agent selection on opinion dynamics in asynchronous BCMS. In the present paper, we study a BCM with heterogeneous agent-selection probabilities, which we implement using node weights. In general terms, we are studying a dynamical process on node-weighted networks. We use node weights to model agents with different probabilities of interacting. These probabilities can encode heterogeneities in individual behavior, such as in sociability or activity levels. We conduct a methodical investigation of the effects of incorporating heterogeneous node weights, which we draw from various distributions, into our generalization of the DW model. We examine these effects on a variety of networks. In our study, we consider fixed node weights that we assign in a way that disregards network structure and agent opinions. However, one can readily adapt the node weights in our BCM to consider a variety of sociological scenarios in which agents have heterogeneous selection probabilities. We find that introducing heterogeneous node weights into our node-weighted BCM results in longer convergence times and more opinion fragmentation than selecting nodes uniformly at random. Moreover, when studying models with heterogeneous node selection, our results illustrate that it is important to consider the baseline influence of assigning node-selection probabilities uniformly at random before drawing conclusions about more specific mechanisms such as algorithmic bias [24]. More generally, our model illustrates the relevance of incorporating node weights into network analysis and dynamics.

Our paper proceeds as follows. In Sec. II, we describe the standard DW model and present our generalized DW model with node weights to incorporate heterogeneous agent-selection probabilities. In Sec. III, we discuss the setup of our simulations, the networks and node-weight distributions that we examine, and the quantities that we compute to characterize the behavior of our model. In Sec. IV, we discuss the results of our numerical simulations of our BCM. In Sec. V, we summarize our results and discuss their implications, present some ideas for future work, and highlight the importance of studying networks with node weights. Our code is available at [46].

II. OUR MODEL

In this section, we first discuss the Deffuant-Weisbuch (DW) [21] bounded-confidence model (BCM) of opinion dy-

namics, and we then introduce our BCM with heterogeneous node-selection probabilities.

A. The standard Deffuant-Weisbuch (DW) BCM

The DW model was introduced over two decades ago [21], and this model and its extensions have been studied extensively since then [4,5]. The DW model was examined originally on complete graphs and encoded agent opinions as scalar values in a closed interval on the real line. Deffuant *et al.* [21] let each agent have an opinion in $[0,1]$, and we follow this convention. The standard DW model has two parameters. The “confidence bound” $c \in [0, 1]$ is a thresholding parameter; when two agents interact, they compromise their opinions by some amount if and only if their opinions differ by less than c . The “compromise parameter” $m \in (0, 0.5]$ (which is also sometimes called a convergence parameter [21] or a cautiousness parameter [26]) parametrizes the amount that an agent changes its opinion as a result of an opinion compromise.

In the standard DW model, agents update their opinions asynchronously. One endows each agent with an initial opinion, which one selects uniformly at random from the interval $[0,1]$. At each discrete time, one uniformly randomly selects a pair of agents to interact. At time t , suppose that one picks agents i and j , whose associated opinions are x_i and x_j , respectively. Agents i and j update their opinions through the following equations:

$$\begin{aligned} x_i(t+1) &= \begin{cases} x_i(t) + m\Delta_{ij}, & \text{if } |\Delta_{ij}(t)| < c \\ x_i(t), & \text{otherwise,} \end{cases} \\ x_j(t+1) &= \begin{cases} x_j(t) + m\Delta_{ij}, & \text{if } |\Delta_{ij}(t)| < c \\ x_j(t), & \text{otherwise,} \end{cases} \end{aligned} \quad (1)$$

where $\Delta_{ij}(t) = x_i(t) - x_j(t)$. When $|\Delta_{ij}(t)| < c$, we say that agents i and j are “receptive” to each other at time t . When $|\Delta_{ij}(t)| \geq c$, we say that agents i and j are “unreceptive” to each other.

When one extends the DW model to consider an underlying network of agents [47], only adjacent agents are allowed to interact. Consider an undirected network $G = (V, E)$, where V is a set of nodes and E is a set of edges between them. Let $N = |V|$ denote the size (i.e., the number of nodes) of a network. Each node of a network represents an agent, and each edge between two agents encodes a social or communication tie between them. At each discrete time, one selects an edge uniformly at random and the two agents that are attached to that edge interact with each other; they update their opinions using Eq. (1).

In the DW model, an alternative to an edge-based approach of randomly selecting an interacting edge is to take a node-based approach to determine the agents that interact. (See Ref. [48] for a discussion of node-based updates versus edge-based updates in the context of voter models.) In a node-based approach, one first randomly selects one node and then randomly selects one of its neighbors. To capture the fact that some agents have more frequent interactions (such as from greater sociability or a stronger desire to share their opinions) than others, we implement a node-based agent-selection procedure in our study.

The choice between edge-based and node-based agent selection can have substantial effects on the dynamics of voter models of opinion dynamics [48], and we expect that this is also true for other types of opinion-dynamics models. We are not aware of existing research that compares edge-based and node-based agent selection in asynchronous BCMs (and, in particular, in DW models), and it seems both interesting and relevant to explore this issue. Most past research on the DW model has considered edge-based selection [5]. However, Refs. [22,24,25] used a node-based selection procedure to model heterogeneous activities of agents.

B. A BCM with heterogeneous node-selection probabilities

We now introduce our BCM with heterogeneous node-selection probabilities. Consider an undirected network $G = (V, E)$. As in the standard DW model, suppose that each agent i has a time-dependent opinion $x_i(t)$. In our BCM, each agent also has a fixed node weight w_i that encodes sociability, how frequently it engages in conversations, or simply the desire to share its opinions. One can think of a node's weight as a quantification of how frequently it communicates with its friends or posts on social media. By incorporating network structure, the standard DW model can include agents with different numbers of friends (or other social connections). However, selecting interacting node pairs uniformly at random is unable to capture the heterogeneous interaction frequencies of individuals. By introducing node weights, we encode such heterogeneity and then examine how it affects opinion dynamics in a BCM. Although we employ fixed node weights, one can adapt our model to include time-dependent node weights, such as through purposeful strategies (e.g., posting more frequently on social media as one's opinions become more extreme).

In our node-weighted BCM, at each discrete time, we first select an agent i with a probability that is proportional to its weight. Agent i then interacts with a neighbor j , which we select with a probability that is equal to its weight divided by the sum of the weights of i 's neighbors. That is, the probabilities

of first selecting agent i and then selecting agent j are

$$P_1(i) = \frac{w_i}{\sum_{k=1}^N w_k}, \quad P_2(j|i) = \frac{w_j}{\sum_{k \in \mathcal{N}(i)} w_k}, \quad (2)$$

where $\mathcal{N}(i)$ denotes the neighborhood (i.e., the set of neighbors) of i and $j \in \mathcal{N}(i)$. Once we select the pair of interacting agents, we update their opinions following the DW opinion update rule in Eq. (1).

Our BCM incorporates heterogeneous node-selection probabilities with node weights that model phenomena such as the heterogeneous sociability of individuals. One can also study heterogeneous selection probabilities of pairwise (i.e., dyadic) interactions, instead of focusing on the probabilities of selecting individuals. For instance, an individual may discuss their ideological views with a close friend more frequently than with a work colleague. One can use edge weights to determine the probabilities of selecting the dyadic interactions in a BCM. One can also relate edge selection to node selection. For example, one can select the edge between nodes i and j either by selecting node i and then node j or by selecting node j and then node i .

III. SIMULATION DETAILS AND CHARACTERIZATION OF OPINIONS

In this section, we discuss the network structures and node-weight distributions that we consider, the setup of our numerical simulations, and the quantities that we compute to characterize steady-state opinions in our simulations.

A. Network structures

We now describe the details of the networks on which we simulate our node-weighted BCM. We summarize these networks in Table I.

We simulate our BCM on complete graphs as a baseline scenario that lets us examine how incorporating

TABLE I. The networks on which we simulate our node-weighted BCM.

Network	Description	Parameters
$C(N)$	Complete graph with N nodes	$N \in \{10, 20, 30, 45, 65, 100, 150, 200, 300, \dots, 1000\}$
$G(N, p)$	Erdős-Rényi (ER) random-graph model with N nodes and homogeneous, independent edge probability p . We let $N = 500$.	$p \in \{0.1, 0.3, 0.5, 0.7\}$
Two-community SBM ^a	Stochastic block model with 2×2 blocks. Edges between nodes in the same set (A or B) exist with a larger probability than edges between nodes in different sets; the block probabilities satisfy $P_{BB} > P_{AA} > P_{AB}$.	$P_{AA} = 49.9/374$ $P_{BB} = 49.9/124$ $P_{AB} = 1/500$
Core-periphery SBM ^a	Stochastic block model with 2×2 blocks. Set A is a set of core nodes and set B is a set of peripheral nodes. The block probabilities satisfy $P_{AA} > P_{AB} > P_{BB}$.	$P_{AA} = 147.9/374$ $P_{BB} = 1/174$ $P_{AB} = 1/25$
Caltech network	The largest connected component of the Facebook friendship network at Caltech on one day in fall 2005. This network, which is part of the FACEBOOK100 data set [49,50], has 762 nodes and 16 651 edges.	

^aOur SBM networks have $N = 500$ nodes. We partition an SBM network into two sets of nodes; set A has 75% of the nodes, and set B has 25% of the nodes.

heterogeneous node-selection probabilities affects opinion dynamics. Although DW models were introduced more than 20 years ago, complete graphs are still the most common type of network that researchers employ in studies of them [4]. To examine finite-size effects, we consider complete graphs of sizes $N \in \{10, 20, 30, 45, 65, 100, 150, 200, 300, \dots, 1000\}$. For all other synthetic networks, we consider networks with $N = 500$ nodes.

To consider networks with different edge densities, we generate synthetic networks using the $G(N, p)$ Erdős-Rényi (ER) random-graph model, where p is the homogeneous, independent probability of an edge between each pair of nodes [51]. When $p = 1$, this yields a complete graph. We examine $G(500, p)$ graphs with $p \in \{0.1, 0.3, 0.5, 0.7\}$.

To determine how a network with an underlying block structure affects the dynamics of our node-weighted BCM, we consider stochastic-block-model (SBM) networks [51] with 2×2 blocks, where each block consists of an ER graph. Inspired by the choices of Kureh and Porter [48], we consider two types of SBM networks. The first SBM has a two-community structure, in which there is a larger probability of edges within a community than between communities. The second SBM has a core-periphery structure, with a set of core nodes with a large probability of edges within the set, a set of peripheral nodes with a small probability of edges within the set, and an intermediate probability of edges between core nodes and peripheral nodes. To construct our 2×2 SBMs, we partition a network into two sets of nodes. Set A has 375 nodes (i.e., 75% of the network) and set B has 125 nodes (i.e., 25% of the network). We define a symmetric edge-probability matrix

$$P = \begin{bmatrix} P_{AA} & P_{AB} \\ P_{AB} & P_{BB} \end{bmatrix}, \quad (3)$$

where P_{AA} and P_{BB} are the probabilities that an edge exists between two nodes in set A and set B, respectively, and P_{AB} is the probability that an edge exists between a node in set A and a node in set B.

In a two-community SBM, the probabilities P_{AA} and P_{BB} are larger than P_{AB} , so edges between nodes in the same community exist with a larger probability than edges between nodes in different communities. For our two-community SBM, we choose P_{AA} and P_{BB} so that the expected mean degree matches that of the $G(500, 0.1)$ ER model if we consider only edges within set A or only edges within set B. A network from the $G(N, p)$ model has an expected mean degree of $p(N - 1)$ [51], so we want the two communities of these SBM networks to have an expected mean degree of $49.9 = 0.1 \times 499$. We thus use the edge probabilities $P_{AA} = 49.9/374$ and $P_{BB} = 49.9/124$. To ensure that there are few edges between the sets A and B, we choose $P_{AB} = 1/500$.

We want our core-periphery SBM with core set A and periphery set B to satisfy $P_{AA} > P_{AB} > P_{BB}$. We chose P_{AA} so that the expected mean degree matches that of the $G(500, 0.3)$ model (i.e., it is 147.9) if we consider only edges within set A. We thus choose the edge probability $P_{AA} = 147.9/374$. To satisfy $P_{AA} > P_{AB} > P_{BB}$, we choose $P_{AB} = 1/25$ and $P_{BB} = 1/174$.

Finally, we investigate our node-weighted BCM on a real social network from Facebook friendship data. We use the Caltech network from the FACEBOOK100 data set [49,50]. Its nodes encode individuals at Caltech, and its edges encode Facebook “friendships” between them on one day in fall 2005. We only consider the network’s largest connected component, which has 762 nodes and 16 651 edges.

B. Node-weight distributions

In Table II, we give the parameters and probability density functions of the node-weight distributions that we examine in our BCM. In this subsection, we discuss our choices of distributions.

To study the effects of incorporating node weights in our BCM, we compare our model to a baseline DW model. To ensure a fair comparison, we implement a baseline DW model that selects interacting agents uniformly at random using a node-based selection process. As we discussed in Sec. I, it is much more common to employ an edge-based selection process. We refer to the case in which all node weights are equal to 1 (that is, $w_i = 1$ for all nodes i) as the “constant weight distribution”. The constant weight distribution (and any other situation in which all node weights equal the same positive number) results in a uniformly random selection of nodes for interaction. This is what we call the “baseline DW model”; we compare our DW model with heterogeneous node weights to this baseline model. We reserve the term “standard DW model” for the DW model with uniformly random edge-based selection of agents. When all of the nodes of a network have the same degree, our baseline DW model is equivalent to the standard DW model.

The node weights in our BCM encode heterogeneities in interaction frequencies, such as when posting content online. The majority of online content arises from a minority of user accounts [52]. A “90-9-1 rule” has been proposed for such participation inequality. In this rule of thumb, about 1% of the individuals in online discussions (e.g., on social-media platforms) account for most contributions, about 9% of the individuals contribute on occasion, and the remaining 90% of the individuals are present online (e.g., they consume content) but do not contribute to it [53]. Participation inequality has been documented in a variety of situations, including in the numbers of posts on digital-health social networks [54], posts on internet support groups [55], and contributions to open-source software-development platforms [56]. Inequality in user activity has also been examined on Twitter [57]. For example, Xiong and Liu [58] used a power-law distribution to model the number of tweets about different topics. A few years ago, a survey by the Pew Research Center found that about 10% of the accounts of adult Twitter users in the United States generate about 80% of the tweets of such accounts [59].

One can interpret the node weights in our BCM as encoding the participation frequencies of individuals who contribute content to a social-media platform. We model online participation inequality by using a Pareto distribution for the node weights. This choice of distribution is convenient because of its simple power-law form. It has also been used to model inequality in a variety of other contexts, including distributions of wealth, word frequencies, website visits, and numbers

TABLE II. The names and specifications of our node-weight distributions. We show both the general mathematical expressions for the means and the specific values of the means for our parameter values. For the Pareto distributions, we truncate the distribution means to four digits after the decimal point. For all other distributions, the means are exact.

Distribution	Probability density function	Parameter values	Domain	Mean	
Constant	$\delta(x - 1)$	N/A	{1}	1	
Pareto-80-10	$\frac{\alpha}{x^{\alpha+1}}$	$\alpha = \log_{4.5}(10)$	$[1, \infty)$	$\frac{\alpha}{\alpha - 1}$	
Pareto-80-20		$\alpha = \log_4(5)$			2.8836
Pareto-90-10		$\alpha = \log_9(10)$			7.2126
Exp-80-10	$\frac{1}{\beta} \exp\left(\frac{-(x-1)}{\beta}\right)$	$\beta = 1.8836$	$[1, \infty)$	$\beta + 1$	
Exp-80-20		$\beta = 6.2125$			21.8543
Exp-90-10		$\beta = 20.8543$			2.8836
Unif-80-10	$\frac{1}{b-1}$	$b = 4.7672$	$[1, b]$	$\frac{1}{2}(1 + b)$	
Unif-80-20		$b = 13.4250$			7.2125
Unif-90-10		$b = 42.7086$			21.8543

of paper citations [60]. When representing social-media interactions, we care only about accounts that make posts or comments; we ignore inactive accounts. Therefore, we impose a minimum node weight in our model. We use the Pareto type-I distribution, which is defined on $[1, \infty)$, so each node has a minimum weight of 1. This positive minimum weight yields reasonable computation times in our simulations of our BCM. Nodes with weights near 0 would have very small probabilities of interacting, and allowing such weights would prolong simulations.

Let Pareto- X - Y denote the continuous Pareto distribution in which (in theory) $X\%$ of the total node weight is distributed among $Y\%$ of the nodes. In practice, once we determine the N node weights for a simulation with a Pareto node-weight distribution, it is not true that precisely $X\%$ of the total weight is held by $Y\%$ of the N nodes. Inspired by the results of the aforementioned Pew Research Center survey of Twitter users [59], we first consider a Pareto-80-10 distribution, in which we expect 80% of the total weight to be distributed among 10% of nodes. The Pareto principle (which is also known as the “80-20 rule”) is a popular rule of thumb that suggests that 20% of individuals have 80% of the available wealth [60]. Accordingly, we also consider a Pareto-80-20 distribution. Finally, as an example of a node-weight distribution with a more extreme inequality, we also consider a Pareto-90-10 distribution.

We also examine uniform and exponential distributions of node weights. To match the domain of our Pareto distributions, we shift the uniform and exponential distributions so that their minimum node weight is also 1. We also choose their parameters to approximate the means of our Pareto distributions. We use Exp- X - Y and Unif- X - Y as shorthand notation to denote exponential and uniform distributions, respectively, with means that match that of the Pareto- X - Y distribution to four decimal places (see Table II). When we examine the results of our numerical simulations, we want to compare distributions with similar means. We use the phrase “80-20 distributions” to refer to the Pareto-80-20, Exp-80-20,

and Unif-80-20 distributions. We analogously use the phrases “80-10 distributions” and “90-10 distributions”. In total, we examine three different families of distributions (Pareto, exponential, and uniform) with tails of different heaviness. In Table II, we show the details of the probability density functions and the parameters of our node-weight distributions.

C. Simulation specifications

In our node-weighted BCM, agents have opinions in the one-dimensional (1D) opinion space $[0,1]$. Accordingly, we suppose that the confidence bound $c \in (0, 1)$ [61]. We suppose that the compromise parameter $m \in (0, 0.5]$, which is the typically studied range for the DW model [4,26]. When $m = 0.5$, two interacting agents that influence each other fully compromise and average their opinions. When $m < 0.5$, the two agents move towards each other’s opinions, but they do not change their opinions to the mean opinion (i.e., they do not fully compromise).

In our node-weighted BCM, the generation of the graphs in a random-graph ensemble, the sets of node weights, the sets of initial opinions, and the selection of pairs of agents to interact at each time step are all stochastic. We use Monte Carlo simulations to reduce these sources of noise in our simulation results. For each of our random-graph models (i.e., the ER and SBM graphs), we generate five graphs. For each graph and each node-weight distribution, we randomly generate ten sets of node weights. For each set of node weights, we generate ten sets of initial opinions that are distributed uniformly at random. In total, we consider 100 distinct sets of initial opinions and node weights for the Monte Carlo simulations of each individual graph. When we compare simulations from different distributions of node weights for the same individual graph, we reuse the same 100 sets of initial opinions.

In theory, the standard DW model and our node-weighted DW model can take infinitely long to approach a steady state. We define an “opinion cluster” S_i to be a maximal connected set of agents in which the pairwise differences in opinions

are all strictly less than the confidence bound c ; adding any other agent to S_r will yield at least one pair of adjacent agents with an opinion difference of at least c . Equivalently, for each graph G , we define the “effective-receptivity network” $G_{\text{eff}}(t) = (V, E_{\text{eff}}(t))$ as the time-dependent subgraph of G with edges only between nodes that are receptive to each other. This set of edges is

$$E_{\text{eff}}(t) = \{(i, j) \in E : |x_i(t) - x_j(t)| < c\}. \quad (4)$$

The opinion clusters are the connected components of the effective-receptivity network $G_{\text{eff}}(t)$. If every pair of opinion clusters, S_a and S_b , is separated by a distance of at least c (i.e., $|x_i - x_j| \geq c$ for all $i \in S_a$ and all $j \in S_b$) at some time \tilde{T} , then (because c is fixed) no agent in S_a can influence the opinion of an agent in S_b (and vice versa) for all $t \geq \tilde{T}$. Meanwhile, agents in each opinion cluster continue to compromise their opinions with each other. Therefore, in finite time, we observe the formation of steady-state clusters of distinct opinions. Inspired by Meng *et al.* [26], we specify that one of our simulations has “converged” if all opinion clusters are separated from each other by a distance of at least c and each opinion cluster has an opinion spread that is less than a tolerance of 0.02. That is, for each cluster S_r , we have $\max_{i,j \in S_r} |x_i - x_j| < 0.02$. We use T to denote the convergence time (which we report as a number of time steps) of a simulation. For each simulation, the connected components of $G_{\text{eff}}(T)$ are the steady-state opinion clusters.

It is computationally expensive to numerically simulate a DW model. Additionally, as we will show in Sec. IV, our node-weighted DW model with heterogeneous node weights often converges to a steady state even more slowly than the baseline DW model. To reduce the computational burden of checking for convergence, we compute the convergence time to three significant digits. We thereby avoid checking for convergence at each time step. To guarantee that each simulation stops in a reasonable amount of time, we set a bailout time of 10^9 time steps. In our simulations, the convergence time is always shorter than the bailout time. We thus report the results of our simulations as steady-state results.

D. Quantifying opinion consensus and fragmentation

In our numerical simulations, we investigate which situations yield consensus (specifically, they result in one “major” opinion cluster, which we will characterize shortly) at steady state and which situations yield opinion fragmentation (when there are at least two distinct major clusters) at steady state [62]. We are also interested in how long it takes to reach steady-state behavior and in quantifying opinion fragmentation when it occurs. To investigate these model behaviors, we compute the convergence time and the number of steady-state opinion clusters. It is common to study these quantities in investigations of BCMs [4,6,26].

In some situations, an opinion cluster has very few agents. Consider a 500-node network in which 499 agents eventually have the same opinion, but the remaining agent (say, Agent 86, despite repeated attempts by Agent 99 and other agents to convince him) retains a distinct opinion at steady state. In applications, it is not appropriate to think of this situation as opinion fragmentation. To handle such situations, we use a




FIG. 1. Sample trajectories of agent opinions versus time t in a single simulation of our node-weighted BCM on a 500-node complete graph. The BCM parameters are $c = 0.1$ and $m = 0.1$, and we draw the node weights from the constant weight distribution. This node-weight distribution gives our baseline DW model. We color the trajectory of each node by its final opinion cluster. Observe that the final opinion clusters have different sizes. There is one minor cluster (in black); it consists of a single node whose final opinion is about 0.4. The major opinion cluster that converges to the largest opinion value has about twice as many nodes as the other major clusters.

notion of “major clusters” and “minor clusters” [34,63]. We characterize major and minor clusters in an *ad hoc* way. We define a “minor” opinion cluster in a network as an opinion cluster with at most 2% of the agents. Any opinion cluster that is not a minor cluster is a “major” cluster. In our simulations, we calculate the numbers of major and minor opinion clusters at steady state. We account only for the number of major clusters when determining if a simulation reaches a consensus state (i.e., exactly one major cluster) or a fragmented state (i.e., more than one major cluster). We still track the number of minor clusters and use the minor clusters when quantifying opinion fragmentation.

Quantifying opinion fragmentation is much less straightforward than determining whether or not there is fragmentation. Researchers have proposed a variety of notions of fragmentation and polarization [64], and they have also proposed several ways to quantify such notions [64–66]. In principle, a larger number of opinion clusters is one indication of more opinion fragmentation. However, as we show in Fig. 1, there can be considerable variation in the sizes (i.e., the number of nodes) of the opinion clusters. For example, suppose that there are two opinion clusters. If the two opinion clusters have the same size, then one can view the opinions in the system as more polarized than if one opinion cluster has a large majority of the nodes and the other opinion cluster has a small minority. Additionally, although we use only major clusters to determine if a system reaches a consensus or a fragmented state, we seek to distinguish quantitatively between scenarios with opinion clusters (including both major and minor clusters) of similar sizes and scenarios with opinion clusters with a large range of sizes. Following Han *et al.* [67], we do this by calculating Shannon entropy.

Suppose that there are K opinion clusters, which we denote by S_r for $r \in \{1, \dots, K\}$. We refer to the set $\{S_r\}_{r=1}^K$ as an “opinion-cluster profile”; such a profile is a partition of a

network. The fraction of agents in opinion cluster S_r is $|S_r|/N$. The Shannon entropy H of the opinion-cluster profile is

$$H = - \sum_{r=1}^K \frac{|S_r|}{N} \ln \left(\frac{|S_r|}{N} \right). \quad (5)$$

The Shannon entropy H gives a scalar value to quantify the distribution of opinion-cluster sizes. For a given opinion-cluster profile, H indicates the increase in information of knowing the opinion-cluster membership of a single agent instead of not knowing the cluster membership of any agent. For a fixed K , the entropy H is larger if the cluster sizes are closer in magnitude than if there is more heterogeneity in the cluster sizes. For opinion-cluster profiles with similar cluster sizes, H is larger if there are more clusters. We use H to quantify opinion fragmentation, with larger H corresponding to more opinion fragmentation. We calculate the steady-state entropy $H(T)$ using all steady-state opinion clusters (i.e., both major and minor clusters).

Another way to quantify opinion fragmentation is to look at a local level and consider the individual agents of a network. As Musco *et al.* [65] pointed out, if an individual agent has many neighbors with similar opinions to it, then it may be “unaware” of other opinions in a network. For example, most of the neighbors of an agent may hold an opinion that is uncommon in the network. This phenomenon is sometimes called a “majority illusion” [68]. If a set of adjacent agents tend to have neighbors with similar opinions as theirs, they may be in an “echo chamber” [69], as it seems that they are largely exposed only to conforming opinions. To quantify the

local observations of the agents of a network, Musco *et al.* [65] examined a notion of local agreement by calculating the fraction of an agent’s neighbors with opinions on the same side of the network’s mean opinion as that agent. In our simulations, we often observe opinion fragmentation with three or more opinion clusters. Therefore, we need to look beyond the mean opinion of an entire network. To do this, we introduce the “local receptiveness” of an agent. At time t , a node i with neighborhood $\mathcal{N}(i)$ has a local receptiveness of

$$L_i(t) = \frac{|\{j \in \mathcal{N}(i) : |x_i(t) - x_j(t)| < c\}|}{|\mathcal{N}(i)|}. \quad (6)$$

That is, $L_i(t)$ is the fraction of the neighbors of agent i at time t to which it is receptive (i.e., with which it will compromise its opinion if they interact). In the present paper, we consider only connected networks, so each agent i has $|\mathcal{N}(i)| \geq 1$ neighbors. If one wants to consider isolated nodes, one can assign them a local receptiveness of 0 or 1. In our numerical simulations, we calculate the local receptiveness of each agent of a network at the convergence time T . We then calculate the mean $\langle L_i(T) \rangle$ of all agents in the network. This is the steady-state mean local receptiveness, as it is based on edges in the steady-state effective-receptivity network $G_{\text{eff}}(T)$. When consensus is not reached, a smaller mean local receptiveness is an indication of more opinion fragmentation. As we will discuss in Sec. IV, computing Shannon entropy and mean local receptiveness can give insight into the extent of opinion fragmentation when one considers them in concert with the number of opinion clusters.

TABLE III. Summary of the trends in our simulations of our node-weighted BCM. Unless we note otherwise, we observe these trends for each of the examined networks (complete graphs, ER and SBM random graphs, and the Caltech Facebook network).

Quantity	Trends
Convergence time	(1) For fixed values of c and m , the heterogeneous node-weight distributions have longer convergence times than the constant weight distribution.
Opinion fragmentation ^a	(1) For fixed values of $c \in [0.1, 0.4]$ and m , the heterogeneous node-weight distributions usually have more opinion fragmentation than the constant weight distribution. (2) For fixed values of c and m and a fixed distribution mean, there usually is more opinion fragmentation when a distribution tail is heavier. (3) For fixed values of c and m and a fixed family of distributions, there usually is more opinion fragmentation when a distribution has a larger mean.
Number of major clusters	(1) A larger minimum value of c is required to always reach consensus for a heterogeneous node-weight distribution than for the constant weight distribution. (2) For fixed values of c and m and a fixed distribution mean, there usually are more major clusters when a distribution tail is heavier. (3) For fixed values of c and m and a fixed family of distributions, there usually are more major clusters when a distribution has a larger mean.
Number of minor clusters	(1) For the constant weight distribution and for fixed c , there usually are more minor clusters when the compromise parameter $m \in \{0.3, 0.5\}$ than when $m = 0.1$. The heterogeneous node-weight distributions do not follow this trend. ^b

^aWe quantify opinion fragmentation using Shannon entropy and mean local receptiveness. We observe clearer trends for Shannon entropy than for mean local receptiveness.

^bFor the Caltech network, we usually observe more minor clusters when $m \in \{0.3, 0.5\}$ than when $m = 0.1$ for each of our heterogeneous weight distributions.

IV. NUMERICAL SIMULATIONS AND RESULTS

In this section, we present results of our numerical simulations of our node-weighted BCM. In our numerical experiments, we consider compromise-parameter values $m \in \{0.1, 0.3, 0.5\}$. For the confidence bound, we first consider $c \in \{0.1, 0.3, 0.5, 0.7, 0.9\}$, and we then examine additional values of c near regions with interesting results. As we discussed in Sec. III C, for each individual graph, we use 100 distinct sets of initial opinions and node weights in Monte Carlo simulations of our BCM. For each of the random-graph models (i.e., ER and SBM graphs), we generate five graphs. For a 500-node complete graph, we use the ten weight distributions in Table II. Because of computation time, we consider the 90-10 distributions only for the 500-node complete graph. For the other networks in Table I, we consider seven distributions: the constant weight distribution, the three 80-10 distributions, and the three 80-20 distributions.

In Table III, we summarize the trends that we observe in the examined networks. In the following subsections, we discuss details of our results for each type of network. The numbers of major and minor clusters, Shannon entropies, and values of mean local receptiveness are all steady-state values. We include our code and figures in our repository at [46]. In the present paper, we visualize our results using heat maps; in our code repository, we also show visualizations with line plots.

A. Simulations on a complete graph

The simplest underlying network structure on which we simulate our node-weighted BCM is a complete graph. A complete graph gives a baseline setting to examine how heterogeneous node-selection probabilities affect opinion dynamics. In our numerical simulations on a 500-node complete graph, we consider all three means (which we denote by 80-10, 80-20, and 90-10) for each of the uniform, exponential, and Pareto node-weight distribution families.

The standard DW model on a complete graph with agents with opinions in the interval $[0,1]$ eventually reaches consensus if the confidence bound $c \geq 0.5$. As one decreases c from 0.5, there are progressively more steady-state opinion clusters (both major and minor clusters) [34,70]. Lorenz [34] showed using numerical simulations that the number of major clusters is approximately $\lfloor \frac{1}{2c} \rfloor$ for the standard DW model. Therefore, there is a transition between consensus and opinion fragmentation for $c \in [0.25, 0.3]$. In our simulations, we observe that this transition occurs for $c \in [0.25, 0.4]$ in our node-weighted BCM. To examine this transition, we thus zoom in on these values of c . For the uniform and exponential distributions, we focus on $c \in [0.25, 0.3]$. For the Pareto distributions, the transition occurs for larger values of c than for the other distributions; we consider additional values of $c \in [0.3, 0.4]$. For the constant weight distribution, which gives our baseline DW model, we examine all values of c that we consider for any other distribution.

In Fig. 2, we show the convergence times of our BCM simulations for various node-weight distributions. For fixed values of c and m , all of the heterogeneous weight distributions yield longer convergence times than the constant weight distribution. Additionally, for fixed c and m and a fixed family of distributions (uniform, exponential, or Pareto), the conver-




FIG. 2. Convergence times (in terms of the number of time steps) in simulations of our node-weighted BCM on a 500-node complete graph. If we consider only the time steps in which interacting nodes actually change their opinions, the convergence times are smaller; however, the trends are the same. For this heat map and all subsequent heat maps, the depicted values are means of simulations of our BCM for each node-weight distribution and each value of the BCM parameter pair (c, m) .

gence time increases as we increase the mean of a distribution. Furthermore, for fixed c and for each heterogeneous weight distribution, the convergence time usually increases as we decrease the compromise parameter m . When calculating convergence time, we include time steps in which two nodes interact but do not change their opinions. To see if the heterogeneous weight distributions have inflated convergence times as a result of having more of these futile interactions, we also calculate the number of time steps to converge when we exclude such time steps. That is, we count the total number of opinion changes that it takes to converge. On a logarithmic




FIG. 3. Steady-state numbers of major opinion clusters in simulations of our node-weighted BCM on a 500-node complete graph with various node-weight distributions. We consider an opinion cluster to be a major cluster if it includes more than 2% of the nodes of a network. (In this case, a major cluster has at least 11 nodes.)

scale, there is little difference between the total number of opinion changes and the total number of time steps that it takes to converge. We include a plot of the numbers of opinion changes in our code repository [46].

In Fig. 3, we show the numbers of steady-state major opinion clusters in our BCM simulations for various node-weight distributions. For all weight distributions, consensus occurs in all of our simulations when the confidence bound $c \geq 0.5$. For fixed values of $c \in [0.1, 0.4]$ and m , the heterogeneous weight distributions usually yield more steady-state major clusters than the constant weight distribution. When we introduce heterogeneous node weights into our BCM, we need a larger confidence bound c than for the constant weight distribution to always reach consensus in our simulations. It appears that our

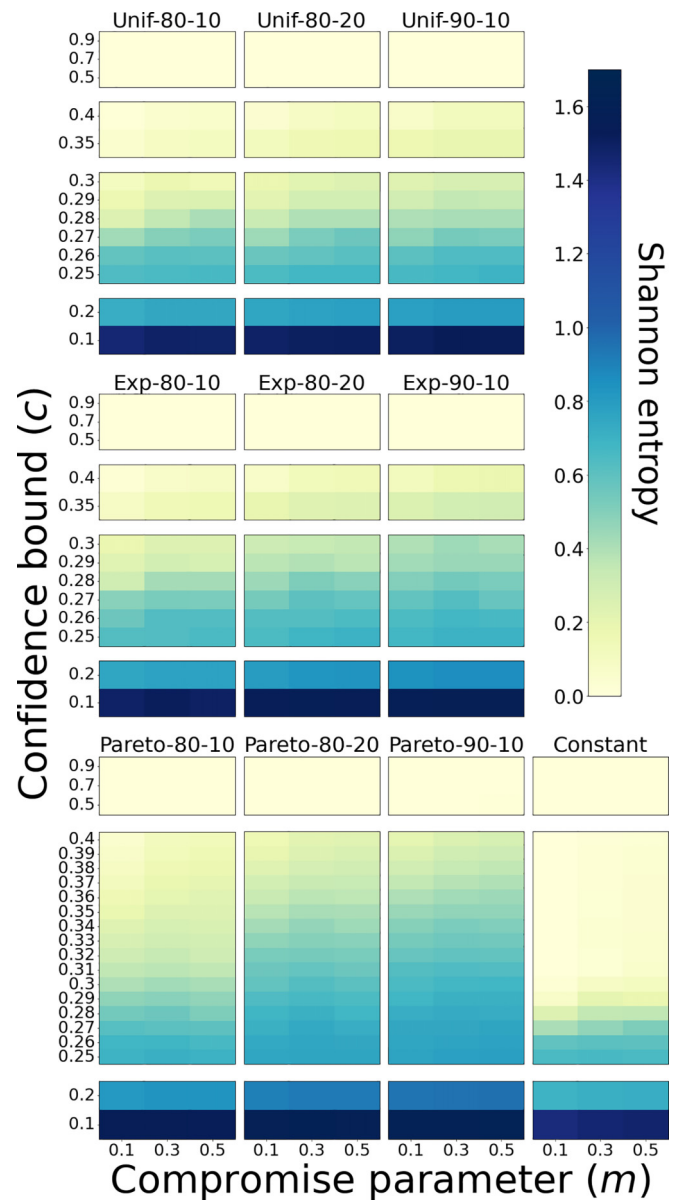


FIG. 4. Shannon entropies of the steady-state opinion-cluster profiles in simulations of our node-weighted BCM on a 500-node complete graph with various node-weight distributions.

BCM with heterogeneous node weights tends to have more opinion fragmentation than the baseline DW model. For fixed c and m , we observe for each distribution family that the mean number of steady-state major clusters increases as we increase the distribution mean. To see this, proceed from left to right in Fig. 3 from the 80-10 distributions to the 80-20 distributions and then to the 90-10 distributions. Additionally, for fixed values of c and m and a fixed distribution mean, there are usually more steady-state major clusters as we proceed from a uniform distribution to an exponential distribution and then to a Pareto distribution.

To investigate how the node-weight distribution and the BCM parameters (i.e., c and m) affect the amount of opinion fragmentation, we calculate the Shannon entropy and mean local receptiveness (see Sec. III D) at steady state. In Fig. 4, we show the steady-state entropies of our BCM simulations

for various node-weight distributions. For all node-weight distributions, when there is opinion fragmentation instead of consensus, the steady-state entropy increases as we decrease the confidence bound c for fixed m . In line with our observations in Fig. 3, when $c \in [0.1, 0.4]$, simulations with heterogeneous weight distributions usually yield larger entropies than simulations with the constant weight distribution. For fixed values of c and m and a fixed distribution mean, we also tend to observe a slightly larger entropy as we proceed from a uniform distribution to an exponential distribution and then to a Pareto distribution. For the Pareto distribution family and fixed values of c and m , the entropy increases as we increase the distribution mean. (Proceed from left to right in Fig. 4.) The exponential and uniform distribution families also have this trend, although it is less pronounced (i.e., the entropies tend to increase only slightly) than for the Pareto distributions. When we quantify fragmentation using Shannon entropy, we conclude that increasing the mean node weight has less effect on the amount of opinion fragmentation for the uniform and exponential distributions than it does for the Pareto distributions. Because Shannon entropy depends on the sizes of the opinion clusters, it provides more information about opinion fragmentation than tracking only the number of major opinion clusters. Our plot of the steady-state mean local receptiveness illustrates the same trends as the entropy. (See our code repository [46] for the relevant figure.) This suggests that both Shannon entropy and mean local receptiveness are useful for quantifying opinion fragmentation.

We now discuss the numbers of steady-state minor opinion clusters in our BCM simulations. (See our code repository [46] for a plot.) For each node-weight distribution and each value of c and m , when we take the mean of our 100 simulations, we obtain at most three steady-state minor clusters. We observe the most minor clusters when $c \in \{0.1, 0.2\}$, which are the smallest confidence bounds that we examine. For the constant weight distribution, we typically observe more minor clusters when $m \in \{0.3, 0.5\}$ than when $m = 0.1$. However, we do not observe this trend for the heterogeneous weight distributions. For example, for the Pareto-80-20 distribution, when $c \in [0.34, 0.4]$, decreasing m results in more minor opinion clusters. For the three Pareto distributions, as we decrease m , we also observe that minor clusters tend to appear at smaller confidence bounds. Smaller values of m entail smaller opinion compromises for interacting agents; this may give more time for agents to interact before they settle into their steady-state opinion clusters. For the constant weight distribution, this may reduce the number of minor clusters by giving more opportunities for agents to assimilate into a major cluster. However, for our heterogeneous weight distributions, nodes with larger weights have larger probabilities of interacting with other nodes and we no longer observe fewer minor clusters as we decrease m .

We now propose a possible mechanism by which our node-weighted BCM may promote the trends in Table III. In Fig. 5, we show the trajectories of opinions versus time for a single simulation with node weights that we draw from a Pareto-80-10 distribution. To qualitatively describe our observations, we examine the large-weight and small-weight nodes (i.e., the nodes that are near and at the extremes of a set of node weights in a given simulation). Because our node-selection probabili-




FIG. 5. Sample trajectories of agent opinions versus time t in a single simulation of our node-weighted BCM on a 500-node complete graph with BCM parameters $c = 0.2$ and $m = 0.1$ and node weights that we draw from a Pareto-80-10 distribution. We color the trajectory of each agent by its node weight, which we normalize so that the sum of all node weights is 1. The nodes in the two minor opinion clusters are all small-weight nodes; their weights are close to 0 (and are hence in purple).

ties are proportional to node weights, we normalize the node weights in a simulation to sum to 1. In Fig. 5, the large-weight nodes appear to quickly stabilize into their associated steady-state major opinion clusters, and some small-weight nodes are left behind to form the two minor clusters. More generally, in our simulations of our BCM on a complete graph, we observe that heterogeneity in the node weights results in large-weight nodes interacting more frequently than other nodes and quickly settling into steady-state major opinion clusters. Small-weight nodes that are not selected for opinion updates early in a simulation are left behind to form the smallest clusters in a steady-state opinion-cluster profile; this increases the amount of opinion fragmentation. When we increase the mean node weight, increase the relative proportion of large-weight nodes (by increasing the heaviness of the tail of a distribution), or decrease the value of the compromise parameter m , small-weight nodes tend to take longer to settle into opinion clusters. Node-weight heterogeneity may thereby promote both opinion fragmentation and the formation of minor opinion clusters.

B. Erdős-Rényi (ER) graphs

We now discuss our simulations of our BCM on $G(N, p)$ ER random graphs, where p is the homogeneous, independent probability of an edge between each pair of nodes [51]. For $p = 1$, these ER graphs are complete graphs. In this subsection, we consider the edge probabilities $p \in \{0.1, 0.3, 0.5, 0.7\}$ and generate five graphs for each value of p . Each graph has $N = 500$ nodes.

For each value of p , we observe the trends in Table III. We include plots of our simulation results for convergence times, steady-state numbers of major and minor opinion clusters, and steady-state values of mean local receptiveness in our code repository [46]. In Fig. 6, we show the steady-state Shannon entropies of our simulations for various node-weight distributions and values of p . The entropies are comparable to those that we obtained in our simulations on a 500-node

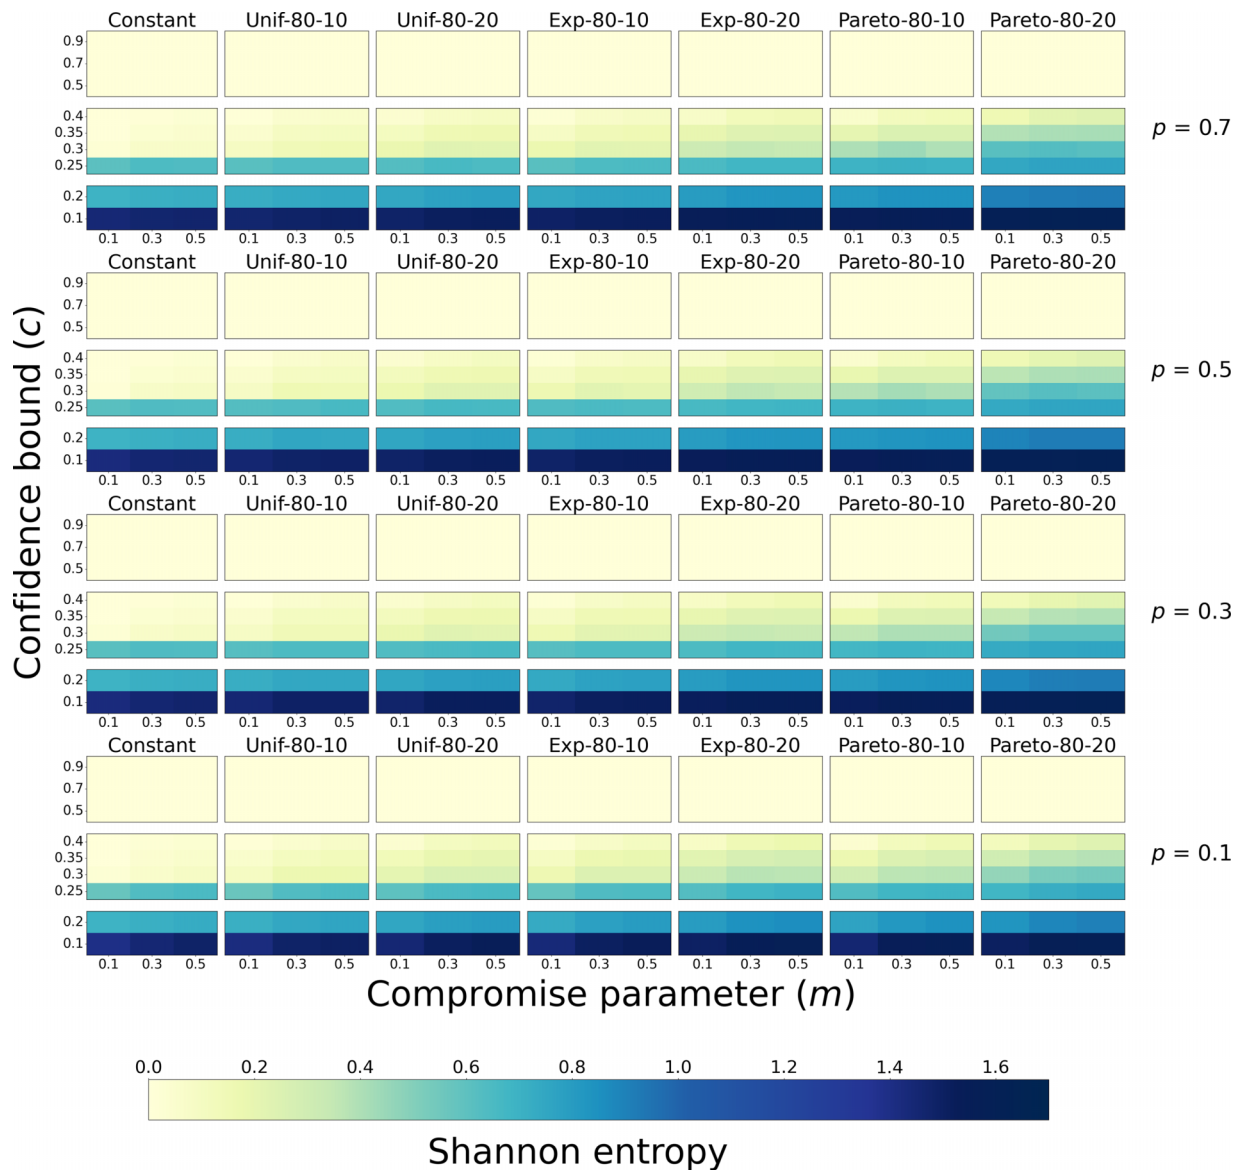


FIG. 6. Shannon entropies of the steady-state opinion-cluster profiles in simulations of our node-weighted BCM on $G(500, p)$ ER random graphs for various node-weight distributions and several values of p .

complete graph (see Sec. IV A). When $c \in [0.1, 0.4]$, for each of our three node-weight distribution families and for fixed values of p , c , and m , the 80-20 distribution tends to yield a larger Shannon entropy than the 80-10 distribution (which has a smaller mean).

For large p , we expect the results of our simulations on $G(500, p)$ graphs to be similar to those for a 500-node complete graph. For $p \in \{0.3, 0.5, 0.7\}$ and $N = 500$, the number of major opinion clusters and the mean local receptiveness are comparable to the corresponding results for a 500-node complete graph. When $p = 0.1$ and there is opinion fragmentation, for a fixed node-weight distribution and fixed values of c and m , we usually observe fewer major opinion clusters than for larger values of p . For $p = 0.1$, a fixed node-weight distribution, and fixed $c \in [0.1, 0.4]$ and m , we also observe that the mean local receptiveness tends to be larger than it is for larger p . One possible contributing factor for this observation may be that smaller values of p yield $G(N, p)$ graphs with more small-

degree nodes; these small-degree nodes have fewer available values of local receptiveness than larger-degree nodes. For example, a node with degree 2 can have a local receptiveness of 0, 0.5, or 1. Unless a small-degree node is an isolated node of the steady-state effective-receptivity network $G_{\text{eff}}(T)$, its presence may help inflate the value of the steady-state mean local receptiveness.

For a fixed node-weight distribution and fixed values of c and m , decreasing p tends to increase the steady-state number of minor opinion clusters. For $p \in \{0.5, 0.7\}$, the steady-state numbers of minor clusters are comparable to the numbers that we obtained for a 500-node complete graph. When $p \in \{0.5, 0.7\}$, for each node-weight distribution and each value of c and m , when we take the mean of our 500 simulations, we obtain at most three steady-state minor clusters. For these simulations, we observe the most minor clusters when $c \in \{0.1, 0.2\}$. For $p = 0.1$, the mean number of steady-state minor clusters is at most nine; this occurs when $c \in \{0.35, 0.4\}$.

It seems sensible that smaller values of p yield more minor opinion clusters. For small p , there are more small-degree nodes than for larger values of p . It is easier for small-degree nodes than for large-degree nodes to be in a minor opinion cluster, as small-degree nodes need to become unreceptive to few neighbors to end up in a minor cluster at steady state. That is, if i is a small-degree node, few neighbors j need to satisfy the inequality $|x_i - x_j| < c$.

C. Stochastic-block-model (SBM) graphs

We now discuss our simulations of our BCM on 500-node SBM random graphs that we generate using the parameters in Table I. For both the two-community and core-periphery SBM graphs, we observe the trends in Table III. We include plots of our simulation results for convergence times, steady-state numbers of major and minor opinion clusters, steady-state Shannon entropies, and steady-state values of mean local receptiveness in our code repository [46].

For the two-community SBM graphs, the steady-state Shannon entropies and numbers of major opinion clusters are comparable to those in our simulations on a 500-node complete graph. When there is opinion fragmentation, for a fixed node-weight distribution and fixed values of c and m , the steady-state values of mean local receptiveness are similar to the values for $G(500, 0.1)$ graphs and tend to be larger than the values for a complete graph. The steady-state numbers of minor opinion clusters are similar to those for the $G(500, 0.1)$ random graphs.

For the two-community SBM graphs, for each node-weight distribution and each value of c and m , when we take the mean of our 500 simulations, we obtain at most ten steady-state minor clusters. We observe the most steady-state minor clusters when $c \in \{0.35, 0.4\}$. Recall that we select the edge probabilities of the two-community SBM so that each of the two communities has an expected mean degree that matches that of $G(500, 0.1)$ graphs. Therefore, it is reasonable that we obtain similar results for the two-community SBM and the $G(500, 0.1)$ random graphs. In our numerical simulations, we assign the node weights without considering the positions (which, in this case, are the community assignments) of the nodes of a network. When we assign weights to nodes uniformly at random, it seems that graph sparsity may be more important than community structure for determining if our BCM reaches a consensus or a fragmented state.

For a fixed node-weight distribution and fixed values of c and m , the core-periphery SBM graphs tend to have fewer steady-state major opinion clusters than a complete graph. Additionally, both the steady-state Shannon entropy and the steady-state mean local receptiveness tend to be larger for the core-periphery SBM graphs than for a complete graph. Larger entropy and smaller mean local receptiveness are both indications of more opinion fragmentation. If we consider only the number of major opinion clusters, it seems that the core-periphery SBM graphs yield less opinion fragmentation than a complete graph. However, when we examine the entire opinion-cluster profile of a network and account for the cluster sizes and the minor clusters, the Shannon entropy reveals that there is more opinion fragmentation for our core-periphery SBM graphs than for a complete graph. The steady-state mean

local receptiveness indicates that the nodes of a core-periphery SBM graph tend to be receptive to a larger fraction of their neighbors than the nodes of a complete graph.

We believe that Shannon entropy gives a more useful quantification of opinion fragmentation than mean local receptiveness. For networks with a large range of degrees, small-degree nodes can inflate the mean local receptiveness. (Analogously, a network's mean local clustering coefficient places more importance than its global clustering coefficient on small-degree nodes [51].) In the context of our node-weighted BCM, consider a node with degree 2 and a node with degree 100, and suppose that both of them have a local receptiveness of 0.5. The larger-degree node's local receptiveness of 0.5 gives a better indication that there may be opinion fragmentation than the smaller-degree node's local receptiveness of 0.5. However, we treat both nodes equally when we calculate the mean local receptiveness. We believe that local receptiveness is a useful quantity to calculate for individual nodes to determine how they perceive the opinions of their neighbors. However, mean local receptiveness appears to be less useful than Shannon entropy for quantifying opinion fragmentation in a network.

For a fixed node-weight distribution and fixed values of c and m , the steady-state numbers of major opinion clusters that we obtain for the core-periphery SBM graphs are comparable to the numbers for a complete graph. The steady-state numbers of minor opinion clusters tend to be larger for core-periphery SBM graphs than for two-community SBM graphs (which have more minor clusters than a complete graph). For each node-weight distribution and each value of c and m , when we take the mean of our 500 simulations, we obtain at most 12 steady-state minor clusters; this occurs when $c = 0.1$. One possibility is that the core-periphery structure makes it easier to disconnect peripheral nodes of an effective-receptivity network, causing these nodes to form minor clusters. For core-periphery SBM graphs, it seems interesting to investigate the effect of assigning node weights in a way that depends on network structure. For example, if we assign all of the large weights to core nodes, will these nodes pull many peripheral nodes into their opinion clusters? If we place a large-weight node in the periphery, will it be able to pull core nodes into its opinion cluster?

D. Caltech network

We now discuss our simulations of our BCM on the Caltech Facebook network, which is an empirical data set in which the nodes represent individuals with Caltech affiliations and the edges represent "friendships" on Facebook on one day in fall 2005 [49,50]. We consider the network's largest connected component, which has 762 nodes and 16 651 edges. The Caltech network has all but one of the trends that we reported in Table III; the only exception is the trend in the number of minor opinion clusters. When there is opinion fragmentation, the Caltech network has more steady-state minor clusters and larger steady-state Shannon entropies than in our synthetic networks.

In Fig. 7, we show the steady-state numbers of minor opinion clusters in simulations of our BCM on the Caltech network. We obtain the most minor clusters when $c = 0.1$, which is the smallest value of c that we examine. Given a

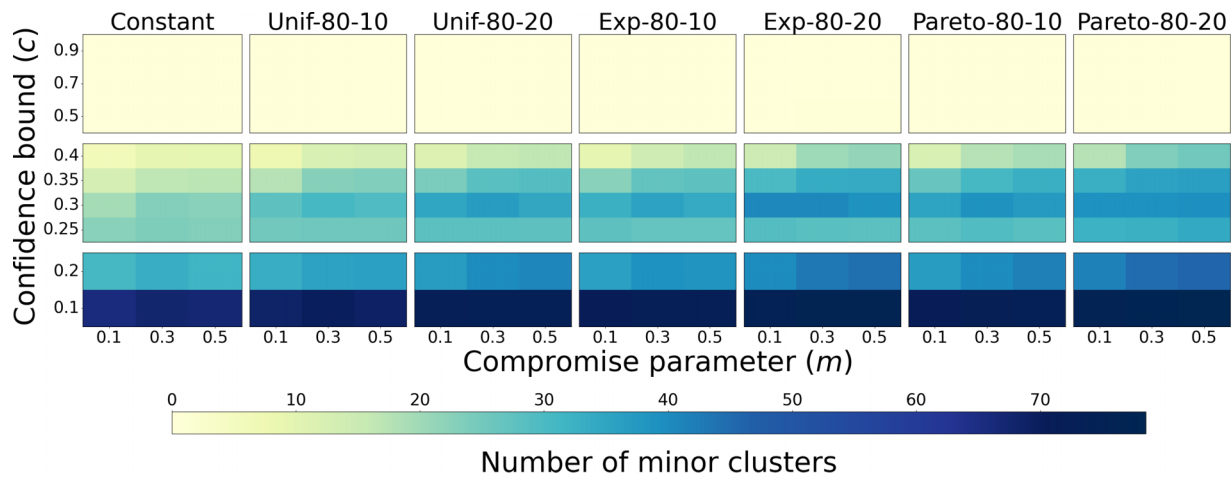


FIG. 7. Steady-state numbers of minor opinion clusters in simulations of our node-weighted BCM on the Caltech Facebook network with various distributions of node weights. We consider an opinion cluster to be a minor cluster if it includes at most 2% of the nodes of a network. (In this case, a minor cluster has at most 15 nodes.)

node-weight distribution and values of c and m , when we take the mean of our 100 simulations on the Caltech network, we often observe a large number of minor clusters (including as many as 78 of them, which is much larger than the single-digit numbers that we usually obtain for our synthetic networks). Additionally, unlike in our synthetic networks, for all node-weight distributions (not just the constant weight distribution), the Caltech network tends to have more minor clusters when $m \in \{0.3, 0.5\}$ than when $m = 0.1$. We include our plot of the steady-state numbers of major opinion clusters in our code repository [46]. For a fixed node-weight distribution and fixed values of c and m , the Caltech network tends to have fewer major opinion clusters than our synthetic networks.

In Fig. 8, we show the steady-state Shannon entropies for the Caltech network. For a fixed node-weight distribution and fixed values of c and m , when there is opinion fragmentation, we observe a larger entropy for the Caltech network than for our synthetic networks. This aligns with our observation that the Caltech network has many more minor opinion clusters than our synthetic networks. We show a plot of the steady-state values of mean local receptiveness for the Caltech

network in our code repository [46]. The values of the mean local receptiveness tend to be larger for the Caltech network than for a 500-node complete graph. We suspect that this arises from the presence of many small-degree nodes in the Caltech network. In Sec. IV C, we discussed the impact of small-degree nodes on the mean local receptiveness.

The histogram of the node degrees of the Caltech network (see Fig. 9) differs dramatically from those of our synthetic networks. Unlike in our synthetic networks, the most common degrees in the Caltech network are among the smallest degrees. In Fig. 9, the tallest bar in the histogram indicates nodes with degrees 1–9. These abundant small-degree nodes are likely to disconnect from large connected components of the effective-receptivity network and form minor opinion clusters. Because we select the initial opinions uniformly at random from $[0,1]$, when $c = 0.1$, it is possible that small-degree nodes are initially isolated nodes of the effective-receptivity network because of their initial opinions. The abundance of small-degree nodes in the Caltech network helps explain its larger steady-state numbers of minor opinion clusters and the correspondingly larger entropies than in our synthetic net-

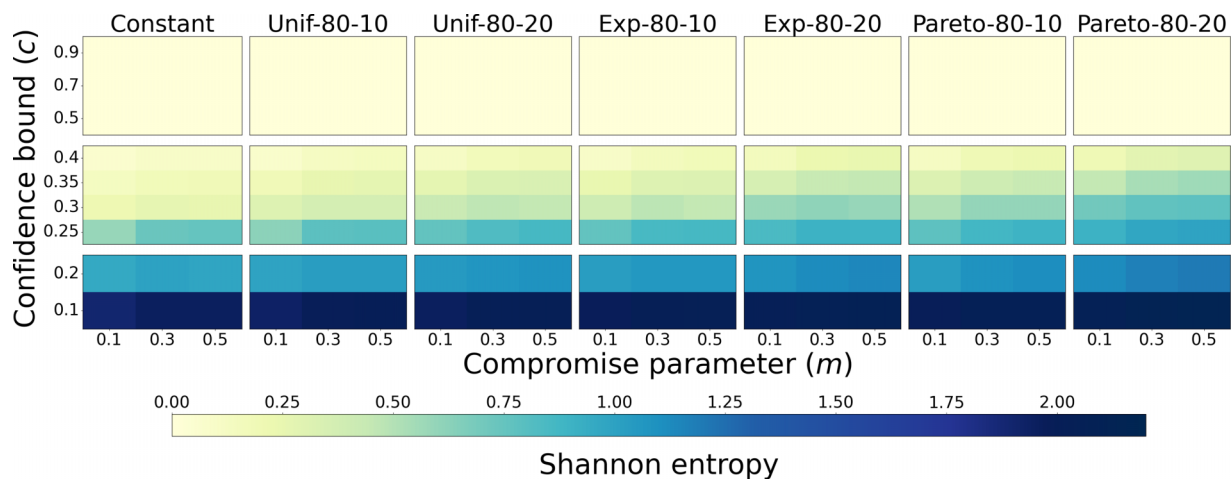


FIG. 8. Shannon entropies of the steady-state opinion-cluster profiles in simulations of our node-weighted BCM on the Caltech Facebook network with various node-weight distributions.




FIG. 9. Histogram of the node degrees of the Caltech Facebook network. The bins have width 10 and originate at the left end point (so the bins indicate degrees of 0–9, 10–19, and so on).

works. Despite the fact that the Caltech network is structurally very different from our synthetic networks, it follows all of the trends in Table III aside from the one for the number of minor opinion clusters. Therefore, it seems that the trends that we observe in our node-weighted BCM when we assign node weights uniformly at random (and hence in a way that is independent of network structure) are fairly robust to the underlying network structure.

E. Finite-size effects

We now investigate finite-size effects by simulating our BCM on complete graphs of different sizes. Previously, to ensure reasonable computation times, we examined synthetic networks with 500 nodes. However, it is useful to get a sense of whether or not the trends in Ta-

ble III hold for networks of different sizes. We thus simulate our BCM on complete graphs of sizes $N \in \{10, 20, 30, 45, 65, 100, 150, 200, 300, \dots, 1000\}$. We examine $m \in \{0.3, 0.5\}$ and $c \in \{0.1, 0.3, 0.5\}$. These values of c give regimes of opinion fragmentation, a transition between fragmentation and consensus for the constant weight distribution, and opinion consensus. We consider the constant weight distribution and the 80-10 distributions (i.e., the uniform, exponential, and Pareto distributions with a mean node weight of 2.8836). We do not examine any larger-mean distributions because they require longer computation times.

We show our results for convergence times and steady-state Shannon entropies. To visualize our results, we plot graph sizes on a logarithmic scale. We include plots of our simulation results at steady state for the numbers of major opinion clusters, the numbers of minor opinion clusters, and the values of mean local receptiveness in our code repository [46].

In Fig. 10, we show the convergence times of our simulations of our BCM on complete graphs of various sizes. For all distributions, the convergence times become longer as we increase the graph size. For each graph size, the convergence times for the heterogeneous weight distributions are similar to each other and are longer than those for the constant weight distribution.

In Fig. 11, we show the steady-state Shannon entropies for our simulations of our BCM on complete graphs of various sizes. For a fixed value of c , we observe similar results when $m = 0.3$ and $m = 0.5$. When $c = 0.5$, for each node-weight distribution, our simulations always reach a consensus (i.e., there is exactly one steady-state major opinion cluster) for $N \geq 200$. Correspondingly, the steady-state entropies are close to 0. (They are not exactly 0 because the calculation of Shannon entropies includes information from minor clusters.) As we increase the network size, the error bars (which indicate one standard deviation from the mean) become progressively

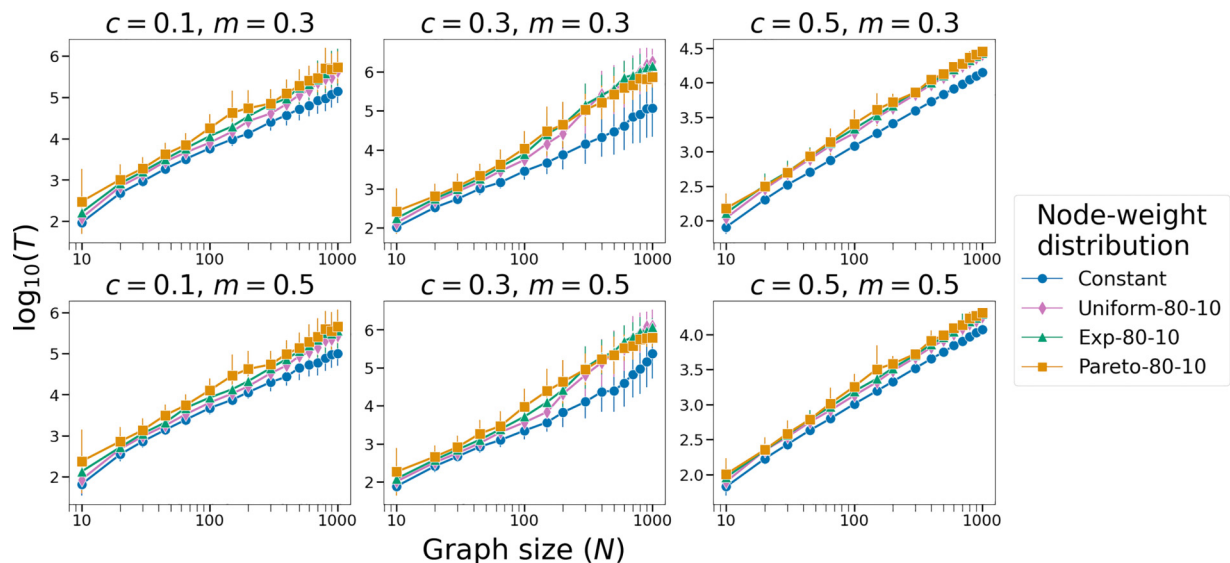


FIG. 10. Convergence times (in terms of the number of time steps) in simulations of our node-weighted BCM on complete graphs of various sizes. We show results for various choices of c and m ; the marker shape and color indicate the node-weight distribution. For this figure and subsequent figures of this type, the points are means of 100 simulations and the error bars indicate one standard deviation from the mean. The horizontal axis gives the graph size on a logarithmic scale. The vertical axes of the plots have different scales.

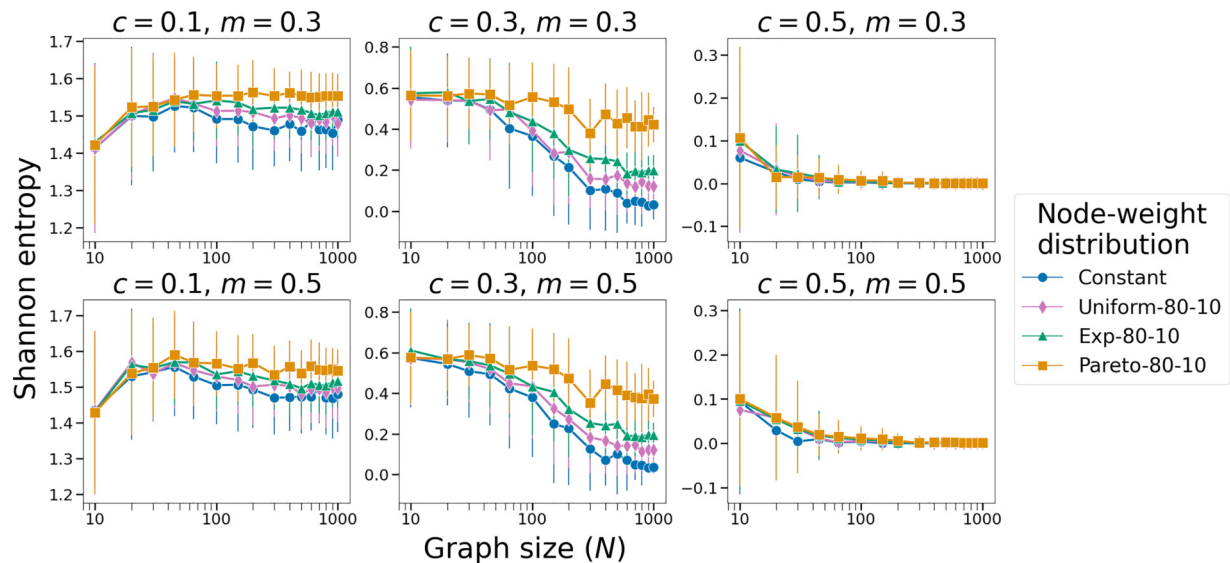


FIG. 11. Shannon entropies of the steady-state opinion-cluster profiles in simulations of our node-weighted BCM on complete graphs of various sizes. We show results for various choices of c and m ; the marker shape and color indicate the node-weight distribution.

smaller. When $c \in \{0.1, 0.3\}$, for sufficiently large graph sizes (specifically, when $N \geq 100$), we observe that the entropy increases as we increase the heaviness of the tail of a distribution. For $c = 0.3$, the mean steady-state entropies appear to no longer change meaningfully with respect to N when $N \geq 400$. For $c = 0.1$, this is the case when $N \geq 100$.

When there is opinion fragmentation, the heterogeneous node-weight distributions tend to yield larger steady-state Shannon entropies (and hence more opinion fragmentation, if one is measuring it using entropy) than the constant weight distribution for each graph size. Additionally, for the 80-10 distributions and sufficiently large graph sizes, we obtain larger entropies as we increase the heaviness of the distribution tail. We have not explored the effect of graph size on the observed trends (see Table III) when we increase the distribution mean for a fixed family of distributions. In our code repository [46], we include a plot that compares the steady-state mean local receptiveness for complete graphs of various sizes. In that plot, we also observe that there tends to be more opinion fragmentation (in the sense of a smaller mean local receptiveness) for heterogeneous node-weight distributions with increasingly heavy tails.

We also examine the steady-state numbers of minor and major opinion clusters in simulations of our BCM on complete graphs of various sizes. We include plots of them in our code repository [46]. For a fixed value of c , we observe similar results when $m = 0.3$ and $m = 0.5$. When $N \leq 49$, there are no minor opinion clusters, by definition, because minor clusters can include at most 2% of the nodes of a network (and even a single node constitutes more than 2% of all nodes for such small networks). When $N \geq 65$ and $c \in \{0.1, 0.3\}$, for each distribution, the number of minor clusters tends to increase as we increase N . We do not observe a clear trend in which node-weight distributions yield more minor clusters. For all N and $c = 0.5$, the mean number of minor clusters is close to 0. We now consider major opinion clusters. When $c = 0.5$ and $N \geq 200$, all simulations yield one major opinion cluster (i.e., they all reach consensus). When $c = 0.3$, for all

graph sizes, there are more major opinion clusters as we increase the heaviness of the tail of a distribution. Additionally, when $c = 0.3$, for the Pareto-80-10 distribution, the number of major clusters tends to increase as we increase the graph size. For the other node-weight distributions, the number of major clusters tends to decrease as we increase the graph size. When $c = 0.1$ and $N \geq 200$, there again tends to be more major clusters as we increase the heaviness of the tail of a distribution, although the trend is not as clear as it was for $c = 0.3$.

Based on our exploration of finite-size effects, we are confident that complete graphs with $N \geq 500$ nodes have the trends in Table III. For graphs with $N \geq 500$ nodes, the mean steady-state Shannon entropies for each node-weight distribution appear to no longer change meaningfully with respect to N . For each graph size, the heterogeneous 80-10 distributions have longer convergence times than the constant weight distribution. For a fixed graph size and fixed values of c and m , we observe more opinion fragmentation as we increase the heaviness of the tail of a distribution. Because of computation time, we have not examined finite-size effects for distributions other than the 80-10 distributions. However, because the mean Shannon entropies no longer change meaningfully with respect to N for graphs with $N \geq 500$ nodes, we hypothesize that the trends in opinion fragmentation and convergence time in Table III continue to hold for our synthetic networks when there are more than 500 nodes.

V. CONCLUSIONS AND DISCUSSION

We studied a bounded-confidence model (BCM) with heterogeneous node-selection probabilities, which we modeled using node weights. One can interpret these node weights as encoding phenomena such as heterogeneous agent sociabilities or activity levels. We studied our node-weighted BCM with fixed node weights that we assign in a way that disregards network structure and node opinions. We demonstrated that our node-weighted BCM has longer convergence times

and more opinion fragmentation than a baseline Deffuant-Weisbuch (DW) BCM in which we uniformly randomly select nodes for interaction. It is straightforward to adapt our BCM to assign node weights in a way that depends on network structure and/or node opinions. See Sec. VB and Sec. VC for discussions.

A. Summary of our main results

We simulated our node-weighted BCM with a variety of node-weight distributions (see Table II) on several random and deterministic networks (see Table I). For each of these distributions and networks, we systematically investigated the convergence time and opinion fragmentation for different values of the confidence bound c and the compromise parameter m . To determine if the nodes of a network reach consensus or if there is opinion fragmentation, we calculated steady-state number of major clusters in our simulations. To quantify the amount of opinion fragmentation, we calculated steady-state Shannon entropy and mean local receptiveness. For a given network, we found that entropy and mean local receptiveness tend to follow the same trends in which node-weight distributions have more opinion fragmentation (see Table III). Additionally, based on our results, we believe that Shannon entropy is more useful than mean local receptiveness for quantifying opinion fragmentation in a network. However, calculating local receptiveness is insightful for explorations of the opinion dynamics of individual nodes.

In our simulations of our node-weighted BCM, we observed a variety of trends (see Table III). In particular, we found that heterogeneous node-weight distributions tend to yield longer convergence times and more opinion fragmentation than the baseline DW model (which we obtain by using a constant weight distribution) in our simulations of our BCM. Opinion fragmentation also tends to increase if either (1) for a fixed distribution mean, we increase the heaviness of the tail of a distribution or (2) for a fixed distribution family, we increase the distribution mean. Given a set of heterogeneous node weights, we hypothesize that large-weight nodes are selected early in a simulation with large probabilities and quickly settle into their associated steady-state major opinion clusters. Small-weight nodes that are not selected early in a simulation are left behind to form small opinion clusters, resulting in more opinion fragmentation than in the baseline DW model.

B. Relating node weights to network structure

We examined deterministic and random graphs with various structures, and we observed the trends in Table III. For each of our BCM simulations, we selected node weights from a specified distribution and then assigned these weights to nodes uniformly at random. Therefore, our investigation conveys what trends to expect with fixed, heterogeneous node weights that are assigned to nodes without regard for network structure. However, our model provides a flexible framework to study the effects of node weights when they are correlated with network structure. For example, one can assign weights to nodes in a way that depends on some centrality measure (such as degree). In our BCM, we expect large-degree and large-weight nodes to have more interactions than small-degree or small-weight nodes. Nodes with larger degrees have more neighbors that can select them for an interaction, and

nodes with larger weights have larger probabilities of being selected for an interaction. One possible area of future work is to investigate the combined effects of node weights and node degrees on the frequencies of interactions and the distribution of steady-state opinions in our BCM. Mean-field approaches, such as the one in [71], may offer insights into these effects.

For a given set of node weights, larger-weight nodes have larger probabilities of interacting with other nodes, so their positions in a network likely influence the dynamics of BCMs and other models of opinion dynamics. One can also investigate the effect of homophily when choosing how to assign node weights. For example, in social-media platforms, very active accounts may engage with each other more frequently by sharing or commenting on each others' posts. We can incorporate such features into our BCM through a positive node-weight assortativity, such that large-weight nodes are more likely to be adjacent to each other than to other nodes.

As in the standard DW model, we assign the initial agent opinions uniformly at random in our BCM. However, in a real social network with community structure, this choice may not be realistic. One can consider a social network with communities with different mean opinion values and examine the effect of placing large-weight nodes in different communities. For example, how does placing all large-weight nodes in the same community affect opinion dynamics and steady-state opinion-cluster profiles? How does the presence of a small community of "outspoken" (i.e., large-weight) nodes influence the final opinions of nodes in other communities of a network? Will the small community quickly induce an echo chamber [69], will it pull other nodes into its opinion cluster, or will something else occur?

C. Relating node weights to node opinions

In the present paper, we considered fixed node weights that are independent of node opinions. One can readily adapt our BCM to incorporate time-dependent node weights, such as ones that depend on node opinions. One can allow the probability of selecting a node for interaction to depend on how extreme its opinion is [22] or on the similarity of its opinion to that of another node [24].

Sîrbu *et al.* [24] studied a modified DW model with heterogeneous node-selection probabilities that model algorithmic bias on social media. In their model, one first selects an agent uniformly at random. One then calculates the magnitude of the opinion difference between that agent and each of its neighbors and then selects a neighbor with a probability that is proportional to this difference. In the context of our BCM, one can represent their agent-selection mechanism using time-dependent node weights. To do this, at each time t , one assigns the same weight to all nodes when selecting a first node i . When selecting a node to interact with i , one then assigns weights to the neighbors j of node i that depend on the opinion difference $|x_i(t) - x_j(t)|$. One assigns a weight of 0 to nodes that are not adjacent to i . The simulations by Sîrbu *et al.* on complete graphs suggest that more algorithmic bias results in longer convergence times and more opinion clusters [24]. Pansanella *et al.* [25] observed similar trends in a study of the algorithmic-bias model of Sîrbu *et al.* for various random-graph models.

In our simulations of our BCM with heterogeneous node-selection probabilities, we observed similar trends of longer convergence times and more opinion clusters (and opinion fragmentation) than in our baseline DW model. When studying a BCM with heterogeneous node-selection probabilities, we see from our results that it is important to consider the baseline influence of assigning node weights uniformly at random before attributing trends such as longer convergence times and more opinion fragmentation to specific mechanisms such as algorithmic bias. Different mechanisms can yield very similar empirical observations.

D. Edge-based heterogeneous activities

In the standard DW model, at each time, one selects an edge of a network uniformly at random and the two agents that are attached to that edge interact with each other [47]. Most past work on the DW model and its extensions has focused on this edge-based selection mechanism [5]. In our BCM, to incorporate node weights (e.g., to encode heterogeneous sociabilities or activity levels of individuals), we instead used a node-based selection mechanism. For voter models of opinion dynamics, it is known that the choice between edge-based and node-based agent selection can substantially affect a model's qualitative behavior [48]. We are not aware of existing research that compares edge-based and node-based agent selection in asynchronous BCMs (and, in particular, in DW models), and it seems interesting to investigate this issue.

Our BCM has node weights to encode heterogeneous activity levels of individuals. One can also examine heterogeneous dyad-activity levels to account for the fact that individuals do not interact with each of their social contacts with the same probability. To encode such heterogeneity, one can construct a variant of our BCM that incorporates edge weights. At each time step, one can select a pair of agents to interact with a probability that is proportional to the weight of the edge between them. Additionally, one can relate edge-selection mechanisms to node-selection mechanisms. We have not examined edge-based heterogeneous activity levels in a BCM, and we expect that it will be interesting to investigate them.

E. Importance of node weights

A key feature of our BCM is our incorporation of node weights into opinion dynamics. Node weights have been used in activity-driven models of temporal networks [39], and activity-driven frameworks have been used to model which agents can interact with each other in models of opinion dynamics [23,40]. In our BCM, the node weights determine the probabilities of selecting agents for interaction in a time-independent network. Alizadeh and Cioffi-Revilla [22], Sîrbu *et al.* [24], and Pansanella *et al.* [25] examined specific scenarios with heterogeneous node-selection probabilities in modified DW models. Our node-weighted BCM gives a general framework to incorporate node weights into asynchronous BCMs. Using our framework, one can consider node weights that are fixed and assigned uniformly at random

(i.e., as we investigated in this paper), are fixed and assigned according to some other probability distribution (see the discussion in Sec. V B), or are assigned in a time-dependent way (see the discussion in Sec. V C).

In network science, node weights have been studied far less than edge weights, and even the term “weighted network” usually refers specifically to edge-weighted networks by default. For example, it is very common to study centralities in edge-weighted networks [72], but studies of centralities in node-weighted networks (e.g., see Refs. [73,74]) are much less common. Heitzig *et al.* [73] generalized common network statistics to node-weighted networks and used node weights to represent the “sizes” of the nodes of a network. They used their framework to study brain networks with node weights that encode the areas of regions of interest, international trade networks with node weights that encode the gross domestic products (GDPs) of countries, and climate networks with node weights that encode areas in a regular grid on the Earth's surface. Singh *et al.* [74] developed centrality measures that account for both edge weights and node weights, and they used them to study service-coverage problems and the spread of contagions. These studies demonstrate the usefulness of node weights for incorporating salient information in network analysis in a variety of applications.

In our node-weighted BCM, we are interested in determining which nodes of a network are (in some sense) more influential than others and thereby have a larger impact on steady-state opinion-cluster profiles. Recently, Brooks and Porter [75] quantified the influence of media nodes in a BCM by examining how their ideologies influence other nodes of a network. An interesting area of future work is to develop ways to quantify the influence of specific nodes in models of opinion dynamics with node weights. For example, can one determine which nodes to seed with extreme opinions to best spread such opinions? Are there nodes that make it particularly easy for communities to reach consensus and remain connected in a steady-state effective-receptivity network $G_{\text{eff}}(T)$? One can tailor the node weights in our BCM to examine a variety of sociological scenarios in which nodes have heterogeneous activity levels or interaction frequencies. More generally, our model illustrates the importance of incorporating node weights into network analysis, and we encourage researchers to spend more time studying the effects of node weights on network structure and dynamics.

ACKNOWLEDGMENTS

We thank A. L. Bertozzi, H. Z. Brooks, J. G. Foster, J. Luo, D. Needell, and the participants of UCLA's Networks Journal Club for helpful comments and discussions. We also thank the two anonymous referees for helpful comments. We acknowledge financial support from the National Science Foundation (Grant No. 1922952) through the Algorithms for Threat Detection (ATD) program. G.J.L. was also supported by National Science Foundation Grant No. 1829071.

[1] C. Castellano, S. Fortunato, and V. Loreto, Statistical physics of social dynamics, *Rev. Mod. Phys.* **81**, 591 (2009).

[2] A. Sîrbu, V. Loreto, V. D. P. Servedio, and F. Tria, Opinion dynamics: Models, extensions and external effects, in

- Participatory Sensing, Opinions and Collective Awareness*, edited by V. Loreto, M. Haklay, A. Hotho, V. D. P. Servodio, G. Stumme, J. Theunis, and F. Tria (Springer International Publishing, Cham, Switzerland, 2017), pp. 363–401.
- [3] S. Lehmann and Y.-Y. Ahn, *Complex Spreading Phenomena in Social Systems: Influence and Contagion in Real-World Social Networks* (Springer International Publishing, Cham, Switzerland, 2018).
 - [4] H. Noorazar, K. R. Vixie, A. Talebanpour, and Y. Hu, From classical to modern opinion dynamics, *Int. J. Mod. Phys. C* **31**, 2050101 (2020).
 - [5] H. Noorazar, Recent advances in opinion propagation dynamics: A 2020 survey, *Eur. Phys. J. Plus* **135**, 521 (2020).
 - [6] A. F. Peralta, J. Kertész, and G. Iñiguez, Opinion dynamics in social networks: From models to data, [arXiv:2201.01322](https://arxiv.org/abs/2201.01322); T. Yasseri (ed.), *Handbook of Computational Social Science* (Edward Elgar Publishing Ltd., Cheltenham, UK, in press).
 - [7] M. Galesic, H. Olsson, J. Dalege, T. van der Does, and D. L. Stein, Integrating social and cognitive aspects of belief dynamics: Towards a unifying framework, *J. R. Soc. Interface* **18**, 20200857 (2021).
 - [8] F. Vazquez, Modeling and analysis of social phenomena: Challenges and possible research directions, *Entropy* **24**, 491 (2022).
 - [9] A. Chacoma and D. H. Zanette, Opinion formation by social influence: From experiments to modeling, *PLoS ONE* **10**, e0140406 (2015).
 - [10] C. Vande Kerckhove, S. Martin, P. Gend, P. J. Rentfrow, J. M. Hendrickx, and V. D. Blondel, Modelling influence and opinion evolution in online collective behaviour, *PLoS ONE* **11**, e0157685 (2016).
 - [11] K. Takács, A. Flache, and M. Mäs, Discrepancy and disliking do not induce negative opinion shifts, *PLoS ONE* **11**, e0157948 (2016).
 - [12] C. Monti, G. De Francisci Morales, and F. Bonchi, Learning opinion dynamics from social traces, in *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD '20* (Association for Computing Machinery, New York, NY, USA, 2020), pp. 764–773.
 - [13] I. V. Kozitsin, Formal models of opinion formation and their application to real data: Evidence from online social networks, *J. Math. Soc.* **46**, 120 (2022).
 - [14] I. V. Kozitsin, Opinion dynamics of online social network users: A micro-level analysis, *J. Math. Soc.* **47**, 1 (2023).
 - [15] D. Carpentras and M. Quayle, Propagation of measurement error in opinion dynamics models: The case of the Deffuant model, *Physica A* **606**, 127993 (2022).
 - [16] M. Mäs, Challenges to simulation validation in the social sciences. a critical rationalist perspective, in *Computer Simulation Validation: Fundamental Concepts, Methodological Frameworks, and Philosophical Perspectives*, edited by C. Beisbart and N. J. Saam (Springer International Publishing, Cham, Switzerland, 2019), pp. 857–879.
 - [17] P. Holme and F. Liljeros, Mechanistic models in computational social science, *Front. Phys.* **3**, 78 (2015).
 - [18] There are also many models of opinion dynamics with discrete-valued opinions and/or polyadic interactions between agents [2,4,76].
 - [19] D. Chandler and R. Munday, *A Dictionary of Media and Communication* (Oxford University Press, Oxford, UK, 2011).
 - [20] R. Hegselmann and U. Krause, Opinion dynamics and bounded confidence: Models, analysis and simulation, *J. Artif. Soc. Social Simul.* **5**, 2 (2002).
 - [21] G. Deffuant, D. Neau, F. Amblard, and G. Weisbuch, Mixing beliefs among interacting agents, *Adv. Complex Syst.* **03**, 87 (2000).
 - [22] M. Alizadeh and C. Cioffi-Revilla, Activation regimes in opinion dynamics: Comparing asynchronous updating schemes, *J. Artif. Soc. Social Simul.* **18**, 8 (2015).
 - [23] J. Zhang, H. Xia, and P. Li, Dynamics of Deffuant model in activity-driven online social network, in *KSS 2018: Knowledge and Systems Sciences, Communications in Computer and Information Science*, Vol. 949, edited by J. Chen, Y. Yamada, M. Ryoike, and X. Tang (Springer Singapore, Singapore, 2018), pp. 215–224.
 - [24] A. Sirbu, D. Pedreschi, F. Giannotti, and J. Kertész, Algorithmic bias amplifies opinion fragmentation and polarization: A bounded confidence model, *PLoS ONE* **14**, e0213246 (2019).
 - [25] V. Pansanella, G. Rossetti, and L. Milli, From mean-field to complex topologies: Network effects on the algorithmic bias model, in *COMPLEX NETWORKS 2021: Complex Networks & Their Applications X*, Studies in Computational Intelligence, Vol. 1016, edited by R. M. Benito, C. Cherifi, H. Cherifi, E. Moro, L. M. Rocha, and M. Sales-Pardo (Springer International Publishing, Cham, Switzerland, 2022), pp. 329–340.
 - [26] 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**, 022312 (2018).
 - [27] A. Hickok, Y. Kureh, H. Z. Brooks, M. Feng, and M. A. Porter, A bounded-confidence model of opinion dynamics on hypergraphs, *SIAM J. Appl. Dyn. Syst.* **21**, 1 (2022).
 - [28] U. Kan, M. Feng, and M. A. Porter, An adaptive bounded-confidence model of opinion dynamics on networks, *J. Complex Netw.* **11**, cnac055 (2023).
 - [29] D. Jacobmeier, Focusing of opinions in the Deffuant model: First impression counts, *Int. J. Mod. Phys. C* **17**, 1801 (2006).
 - [30] A. Carro, R. Toral, and M. San Miguel, The role of noise and initial conditions in the asymptotic solution of a bounded confidence, continuous-opinion model, *J. Stat. Phys.* **151**, 131 (2013).
 - [31] P. Sobkowicz, Extremism without extremists: Deffuant model with emotions, *Front. Phys.* **3**, 17 (2015).
 - [32] G. Weisbuch, G. Deffuant, F. Amblard, and J.-P. Nadal, Meet, discuss, and segregate!, *Complexity* **7**, 55 (2002).
 - [33] G. Deffuant, F. Amblard, and G. Weisbuch, How can extremism prevail? A study based on the relative agreement interaction model, *J. Artif. Soc. Social Simul.* **5**, 1 (2002).
 - [34] J. Lorenz, Continuous opinion dynamics under bounded confidence: A survey, *Int. J. Mod. Phys. C* **18**, 1819 (2007).
 - [35] G. Kou, Y. Zhao, Y. Peng, and Y. Shi, Multi-level opinion dynamics under bounded confidence, *PLoS ONE* **7**, e43507 (2012).
 - [36] J. Zhang, Convergence analysis for asymmetric Deffuant-Weisbuch model, *Kybernetika* **50**, 32 (2014).
 - [37] C. Huang, Q. Dai, W. Han, Y. Feng, H. Cheng, and H. Li, Effects of heterogeneous convergence rate on consensus in opinion dynamics, *Physica A* **499**, 428 (2018).

- [38] G. Chen, W. Su, W. Mei, and F. Bullo, Convergence properties of the heterogeneous Deffuant-Weisbuch model, *Automatica* **114**, 108825 (2020).
- [39] N. Perra, B. Gonçalves, R. Pastor-Satorras, and A. Vespignani, Activity driven modeling of time varying networks, *Sci. Rep.* **2**, 469 (2012).
- [40] D. Li, D. Han, J. Ma, M. Sun, L. Tian, T. Khouw, and H. E. Stanley, Opinion dynamics in activity-driven networks, *Europhys. Lett.* **120**, 28002 (2017).
- [41] N. Masuda, N. Gibert, and S. Redner, Heterogeneous voter models, *Phys. Rev. E* **82**, 010103(R) (2010).
- [42] A. Baronchelli, C. Castellano, and R. Pastor-Satorras, Voter models on weighted networks, *Phys. Rev. E* **83**, 066117 (2011).
- [43] S. Huet, G. Deffuant, and W. Jager, A rejection mechanism in 2D bounded confidence provides more conformity, *Adv. Complex Syst.* **11**, 529 (2008).
- [44] M. McPherson, L. Smith-Lovin, and J. M. Cook, Birds of a feather: Homophily in social networks, *Annu. Rev. Soc.* **27**, 415 (2001).
- [45] D. Spohr, Fake news and ideological polarization: Filter bubbles and selective exposure on social media, *Bus. Inf. Rev.* **34**, 150 (2017).
- [46] See <https://gitlab.com/graceli1/NodeWeightDW>.
- [47] G. Weisbuch, G. Deffuant, F. Amblard, and J. P. Nadal, Interacting agents and continuous opinions dynamics, in *Heterogenous Agents, Interactions and Economic Performance*, edited by R. Cowan and N. Jonard (Springer-Verlag, Heidelberg, Germany, 2003), pp. 225–242.
- [48] Y. H. Kureh and M. A. Porter, Fitting in and breaking up: A nonlinear version of coevolving voter models, *Phys. Rev. E* **101**, 062303 (2020).
- [49] V. Red, E. D. Kelsic, P. J. Mucha, and M. A. Porter, Comparing community structure to characteristics in online collegiate social networks, *SIAM Rev.* **53**, 526 (2011).
- [50] A. L. Traud, P. J. Mucha, and M. A. Porter, Social structure of Facebook networks, *Physica A* **391**, 4165 (2012).
- [51] M. E. J. Newman, *Networks*, 2nd ed. (Oxford University Press, Oxford, UK, 2018).
- [52] L. Guo, E. Tan, S. Chen, X. Zhang, and Y. E. Zhao, Analyzing patterns of user content generation in online social networks, in *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '09 (Association for Computing Machinery, New York, NY, USA, 2009), pp. 369–378.
- [53] J. Nielsen, The 90-9-1 rule for participation inequality in social media and online communities, Nielsen Norman Group, <https://www.nngroup.com/articles/participation-inequality/> (2006).
- [54] T. van Mierlo, The 1% rule in four digital health social networks: An observational study, *J. Med. Internet Res.* **16**, e33 (2014).
- [55] B. Carron-Arthur, J. A. Cunningham, and K. M. Griffiths, Describing the distribution of engagement in an internet support group by post frequency: A comparison of the 90-9-1 principle and Zipf's law, *Internet Interventions* **1**, 165 (2014).
- [56] M. Gasparini, R. Clarisó, M. Brambilla, and J. Cabot, Participation inequality and the 90-9-1 principle in open source, in *Proceedings of the 16th International Symposium on Open Collaboration*, OpenSym '20 (Association for Computing Machinery, New York, NY, USA, 2020), article number 6.
- [57] A. Antelmi, D. Malandrino, and V. Scarano, Characterizing the behavioral evolution of Twitter users and the truth behind the 90-9-1 rule, in *Companion Proceedings of the 2019 World Wide Web Conference*, WWW '19 (Association for Computing Machinery, New York, NY, USA, 2019), pp. 1035–1038.
- [58] F. Xiong and Y. Liu, Opinion formation on social media: An empirical approach, *Chaos* **24**, 013130 (2014).
- [59] S. Wojcik and A. Hughes, Sizing up Twitter users, <https://www.pewresearch.org/internet/2019/04/24/sizing-up-twitter-users/>, Pew Research Center (2019).
- [60] M. E. J. Newman, Power laws, Pareto distributions and Zipf's law, *Contemp. Phys.* **46**, 323 (2005).
- [61] The extreme case $c = 0$ is degenerate (because no agents update their opinions), and the case $c = 1$ allows all adjacent agents to interact with each other. We are not interested in examining these cases.
- [62] Some researchers use the term “polarization” to refer to the presence of exactly two opinion clusters (or of exactly two major opinion clusters) and “fragmentation” to refer to the presence of three or more opinion clusters (or of three or more major opinion clusters) [20,64]. However, because we are interested in distinguishing between consensus states and any state that is not a consensus, we use the term “fragmentation” for any state with at least two major opinion clusters. We then quantify the extent of opinion fragmentation.
- [63] M. F. Laguna, G. Abramson, and D. H. Zanette, Minorities in a model for opinion formation, *Complexity* **9**, 31 (2004).
- [64] A. Bramson, P. Grim, D. J. Singer, S. Fisher, W. Berger, G. Sack, and C. Flocken, Disambiguation of social polarization concepts and measures, *J. Math. Soc.* **40**, 80 (2016).
- [65] C. Musco, I. Ramesh, J. Ugander, and R. T. Witter, How to quantify polarization in models of opinion dynamics, [arXiv:2110.11981](https://arxiv.org/abs/2110.11981).
- [66] J. A. Adams, G. White, and R. P. Araujo, Mathematical measures of societal polarisation, *PLoS ONE* **17**, e0275283 (2022).
- [67] W. Han, Y. Feng, X. Qian, Q. Yang, and C. Huang, Clusters and the entropy in opinion dynamics on complex networks, *Physica A* **559**, 125033 (2020).
- [68] K. Lerman, X. Yan, and X.-Z. Wu, The “majority illusion” in social networks, *PLoS ONE* **11**, e0147617 (2016).
- [69] S. Flaxman, S. Goel, and J. M. Rao, Filter bubbles, echo chambers, and online news consumption, *Public Opin. Q.* **80**, 298 (2016).
- [70] E. Ben-Naim, P. Krapivsky, and S. Redner, Bifurcations and patterns in compromise processes, *Physica D* **183**, 190 (2003).
- [71] S. C. Fennell, K. Burke, M. Quayle, and J. P. Gleeson, Generalized mean-field approximation for the Deffuant opinion dynamics model on networks, *Phys. Rev. E* **103**, 012314 (2021).
- [72] T. Opsahl, F. Agneessens, and J. Skvoretz, Node centrality in weighted networks: Generalizing degree and shortest paths, *Social Netw.* **32**, 245 (2010).

- [73] J. Heitzig, J. F. Donges, Y. Zou, N. Marwan, and J. Kurths, Node-weighted measures for complex networks with spatially embedded, sampled, or differently sized nodes, *Eur. Phys. J. B* **85**, 38 (2012).
- [74] A. Singh, R. R. Singh, and S. R. S. Iyengar, Node-weighted centrality: A new way of centrality hybridization, *Comput. Social Netw.* **7**, 6 (2020).
- [75] H. Z. Brooks and M. A. Porter, A model for the influence of media on the ideology of content in online social networks, *Phys. Rev. Res.* **2**, 023041 (2020).
- [76] F. Battiston, G. Cencetti, I. Iacopini, V. Latora, M. Lucas, A. Patania, J.-G. Young, and G. Petri, Networks beyond pairwise interactions: Structure and dynamics, *Phys. Rep.* **874**, 1 (2020).

Correction: The second sentence after Eq. (4) contained errors and has been fixed. A sentence has been added immediately after the corrected sentence.

Second Correction: The second and third sentences of the fourth paragraph in Sec. I contained errors and have been fixed.