A model for the influence of media on the ideology of content in online social networks

Heather Z. Brooks and Mason A. Porter
Department of Mathematics, University of California Los Angeles, Los Angeles, California 90095, USA

(Received 19 April 2019; revised manuscript received 29 May 2019; accepted 14 January 2020; published 14 April 2020)

Many people rely on online social networks as sources of news and information, and the spread of media content with ideologies across the political spectrum influences online discussions and impacts offline actions. To examine the impact of media in online social networks, we generalize bounded-confidence models of opinion dynamics by incorporating media accounts as influencers in a network. We quantify partisanship of content with a continuous parameter on an interval, and we formulate higher-dimensional generalizations to incorporate content quality and increasingly nuanced political positions. We simulate our model with one and two ideological dimensions, and we use the results of our simulations to quantify the "entrainment" of content from nonmedia accounts to the ideologies of media accounts in a network. We maximize media impact in a social network by tuning the number of media accounts and the numbers of followers of those accounts. Using numerical computations, we find that the entrainment of the ideology of content that is spread by nonmedia accounts to media ideology depends on a network's structural features, including its size, the mean number of followers of its nodes, and the receptiveness of its nodes to different opinions. We then introduce content quality—a key novel contribution of our work—into our model. We incorporate multiple media sources with ideological biases and quality-level estimates that we draw from real media sources and demonstrate that our model can produce distinct communities ("echo chambers") that are polarized in both ideology and quality. Our model provides a step toward understanding content quality and ideology in spreading dynamics, with ramifications for how to mitigate the spread of undesired content and promote the spread of desired content.

DOI: 10.1103/PhysRevResearch.2.023041

I. INTRODUCTION

Online social media (such as Twitter, Facebook, Instagram, and others) have become extremely influential sources of news in daily life. For example, over two thirds of American adults who participated in a recent survey responded that social media, and their associated networks, is their primary source for obtaining news [1]. Given this large audience and the ease of sharing content online, the content that spreads on online social networks can have important consequences on public opinion, policy, and voting [2,3].

The spread of content in a social network depends on the actions and biases of the individuals in that network. Individual user preferences play a strong role in the choice to consume news that conforms to (or even enhances) previously held views [4]. One reason that propaganda, misinformation, and disinformation have become so widespread on social-media platforms is that users are more likely to share a false or misleading story if it seems to confirm or support their biases [5,6]. Another challenge is that content is also spread and amplified through bot, cyborg, and sockpuppet accounts [7].

The political biases of content and the accounts that spread it goes beyond a naive separation of "liberal" versus "conservative," and the political spectrum of ideology can include axes for social views ("progressive" versus "conservative"), economic views ("socialist" versus "capitalist"), views on government involvement ("libertarian" versus "authoritarian"), and others. Opinions can also be issue-dependent. In a large body of research, Poole and Rosenthal [8,9] developed methods to classify ideological positions in roll-call voting in the United States Congress in ideological space, and they found that two dimensions is typically sufficient to capture an overwhelming majority of the variance of ideologies. Their work laid the foundation of the Voteview project [10], which characterizes the ideologies of legislators in the United States Congress based on their voting records. Multidimensional political ideologies have also been explored in the context of the United States Supreme Court. For example, by examining singular value decompositions, Sirovich [11] demonstrated that decisions by the judges in the second Rehnquist Court are well-described by a two-dimensional space. This has typically also been true of voting on legislation in the United States Congress [12].

Political biases are not the only way in which people judge the news that they consume and share on social media. As the prevalence of misinformation, disinformation, and "fake news" becomes increasingly prominent in everyday life and the global conversations of talking heads, social-media users are also considering the quality of news content. For example, in the aforementioned Pew Research Study [1], the

majority of respondents acknowledged that the news that they see, consume, and spread on social-media platforms may be inaccurate. There is evidence that poor-quality content travels deeper, more broadly, and faster than high-quality content (such as fact-based reporting) [13]. The Pew Research study [1] also concluded that these differences in spreading patterns are primarily the result of human actions, rather than arising from bot or cyborg accounts.

A large body of work has examined the spread of content on social media [13,14]. One approach is to mathematically model the spread of content on social media using ordinary differential equations, where the equations can capture changes in time of the proportion of a population that is "susceptible to," "exposed to," "infected with," or "immune to" the propagation of such content (e.g., a rumor) [15–18]. Such compartmental models have the advantage of being analytically tractable, but they do not capture the effects of either network structure or heterogeneity in account characteristics. Other studies have focused instead on data-driven or computational approaches [4,5,13,19]. One important research direction in this vein is quantifying the existence of distinct communities and "echo chambers" [20–23] on social media, wherein individuals interact primarily with other like-minded individuals. Some recent work on modeling the spread of content has tried to bridge the gap between these two approaches by introducing mathematical models that capture network features of social media [24,25], including some very recent mechanistic models of radicalization dynamics [26] and filter bubbles [27]. Our approach in the present paper has some similarities with those in Prasetya and Murata [28], who generalized an independent-cascade model to explore the effects of selective exposure and connection strengths on news propagation, and Martins [29], who developed a continuous-opinion, discrete-action model to explore the emergence of extremism. Extremism in models of opinion dynamics was also studied recently in the context of radicalization and terrorism [30].

In the present paper, we formulate and study a model for the influence of media accounts on the ideology and quality of content that is shared in an online social network. First, we develop a general content-spreading model with an n-dimensional continuous ideology space and spreading choices that are based on a bounded-confidence mechanism [31–33]. We also model how to include media accounts as influencer nodes in a network. This has been considered in other recent work, such as in the context of voters models with discrete-valued opinions [34,35]. In our work, we conduct extensive numerical simulations to examine the impact of media accounts on content in examples with one and two ideological dimensions. We then introduce content quality into our model; this is a key novel feature of our work. Using the resulting enhanced model, we employ numerical simulations to examine the effect of multiple media sources on the "entrainment" (with one possible uncharitable interpretation as "brainwashing") of nonmedia nodes in the network by media nodes. As a parallel to our consideration of media nodes, radical or charismatic leaders have also been introduced into voter models [36,37], DeGroot models [38], and bounded-confidence models [39].

Our work advances the study of online content spreading in several ways. In our model, we propose a measure of media impact; this allows us to quantify the influence that a set of media accounts has on the ideology of content in a social network at consensus. Our model has the advantageous feature of supporting a multidimensional continuous ideology space for content, although we are not the first to employ a multidimensional opinion space. See Refs. [40,41] for a discussion of consensus in bounded-confidence models in $\mathbb{R}^d$ and Ref. [42] for some results on the convergence and stability of such models. For wider-ranging discussions of consensus dynamics and opinion models, see Refs. [31,43]. A key novelty of our work is the introduction of content quality into the spreading dynamics; to our knowledge, our model is the first to explore the effects of content quality on spreading. This is an important advancement for two reasons: (1) the quality of news content has a major effect on its spreading dynamics, and it is therefore important to study it from a dynamical perspective; and (2) the spread of poor-quality news content is a major social problem. Investigations into such problems—especially ones that produce potential mechanistic insights into them—have the potential to help motivate strategies for mitigating their effects.

Our paper proceeds as follows. In Sec. II, which gives a concise overview of our work, we give a nontechnical sketch of our model and briefly outline our key results. In Sec. III, we develop our bounded-confidence model with content spreading. We discuss network structures that incorporate both nonmedia and media accounts in Sec. III A, and we introduce our content updating rule in Sec. III B. In Sec. IV, we examine the dynamics of our model with one ideological dimension for content. We describe how we simulate the model in Sec. IV A, and we quantify media impact for different parameter values and network architectures in Sec. IV B. We present a generalization of our model to two ideological dimensions for content in Sec. V. In Sec. VI, we develop our model further to incorporate a notion of media quality into its content updating rule. In Sec. VII, we show simulations and results of media impact with the incorporation of content quality, including an example in which we draw media biases and qualities from a publicly available, hand-curated distribution. We conclude and discuss several possible extensions of our work in Sec. VIII.

II. OVERVIEW OF OUR MODEL AND RESULTS

The aim of our paper is to advance understanding of the spread of misinformation and "fake news" on social media through mathematical modeling. This contrasts with the more common data-driven, computational approaches to these problems that tend to use existing models (such as compartmental models). An important facet of our approach is that the underlying spreading mechanisms are based on continuous and potentially multidimensional quantities—specifically, ideology and quality of content—that result in discrete actions (spreading a message). By incorporating these continuous variables into the spreading dynamics, our model (1) provides a plausible framework to study the spread of extremism on networks and (2) and lays the foundation for systematic comparison with empirical data of the spread of ideologies online.

Our model takes the form of a dynamical process on a network. It includes (1) a time-independent network structure

of accounts and relationships between them in an online social network (see Sec. III A) and (2) a mechanism for the temporal evolution of the ideology of the content that is produced by media accounts in that network (see Sec. III B). To represent the network structure, we construct a graph in which nodes represent accounts and edges represent followership relationships between these accounts. One can either construct these networks from empirical data (such as Facebook or Twitter networks) or using synthetic networks (e.g., from generative models of random graphs). In a network, we designate each account as either a media account or a nonmedia account. *Media accounts* are accounts that do not follow any other accounts but are followed by at least one other account. In our model, we construe media accounts as "influencers," "opinion sources," or "external opinion forcing" in a network. It is for simplicity that we assume that their content is not influenced by any other accounts. All other accounts in a network are *nonmedia accounts*, which each follow at least one other account in the network.

We model the content that spreads in an online social network based on its ideology (see Sec. VII) and eventually also on its quality (see Sec. VI). We suppose that accounts in a network spread content during each discrete time step. We determine the ideology of the content that is spread by an account at each time step with a simple update rule. A helpful way to conceptualize the update rule is as follows. At each time step, each account "reads" the content from all of the accounts that it follows, and it is influenced only by accounts whose content is sufficiently similar to the ideology of its own content. In the next time step, the account spreads new content whose ideology is a mean of the account's old ideology and the aforementioned sufficiently similar ideologies of the accounts that it reads. As a real-life example, a user may read several articles on a topic and then create a message (e.g., a tweet) based on these articles that has their own "spin" on the material. We give a precise definition of the update rule in Sec. III B.

An important question that we are able to study with our model is how much the ideology of the media account(s) in a network entrains the content that is spread by nonmedia accounts in that network. That is, how strong does the external forcing by media need to be to have a concrete impact on the outcome of the ideological dynamics? To quantify the amount of media influence, we calculate a summary statistic that we call *media impact*. Using numerical simulations of our model, we illustrate how the number of media accounts and the number of followers per media account affect the media impact. Surprisingly, we find that the most successful media entrainment occurs when there are a moderate number of media accounts that each have a moderate number of followers. We study how network parameters (such as the number of nonmedia accounts) affect this phenomenon in Secs. IV and V.

An exciting outcome of our modeling and analysis is that "echo chambers" arise naturally from our content-spreading mechanism. Following common usage in sociological research (see, e.g., Refs. [20,21]), we use the term "echo chamber" to refer to a group of nonmedia nodes that primarily or solely influence each other's ideologies and are not influenced much by accounts that are outside the group. These ideological echo chambers emerge even when there are many followership connections between accounts in the different echo chambers. (In our model, the follower relationships do not change, although it would be fascinating to study a generalization of our model in the form of a coevolving network.) In Sec. VII, we use a hand-curated distribution of ideologies and qualities of real-world media outlets as an input to our model. We observe the emergence of two polarized communities: one with ideologically moderate, high-quality content and one with conservative, low-quality content. Our model is very flexible, as it allows many generalizations and modifications to study features of interest in online social networks. For example, one can study the influence of multiple social-media platforms by generalizing our networks to multilayer networks, and one can also draw media content and ideology from different probability distributions. It is also possible to adapt our model to study other applications in which one combines continuous parameters (for ideology, quality, self-confidence, or something else) with discrete actions based on those parameters. A few examples include models for gambling in sports, choices between competing products, and choices of students of majors and courses to take at a university.

III. A BOUNDED-CONFIDENCE MODEL WITH CONTENT SPREADING

A. Social network structure

Consider a social network with N nonmedia accounts. We represent this social network with a graph $G(\mathcal{V}_N, \mathcal{E})$, where each node in the set $\mathcal{V}_N$ represents an account and each edge in the set $\mathcal{E}$ represents a follower relationship from one account to another. Specifically, we say that account i is a *follower* of account j if there is a directed edge from j to i. With this structure, we can represent this social network as an adjacency matrix A, where $A_{ij} = 1$ if account i follows account j and $A_{ij} = 0$ otherwise. We do not assume reciprocal follower relationships, so A is not symmetric in general. We also assume that the network structure is fixed, so follower relationships do not change in time.

We now introduce media accounts into the above social network. Suppose that media accounts do not change their ideology, as encapsulated in the content that they produce, so they are not influenced by other accounts. Therefore, the only edges in media accounts are ones that are directed outward. In other words, the in-degree of a media account is always 0. Of course, this is a simplification of real-world media outlets, which may (in some cases) be swayed by individuals or by public opinion. We make this simplifying assumption for two reasons: (1) we suppose that the media ideology represents the ideology of a particular topic or news story (and such a view of a given media outlet is unlikely to change much on the timescale of online content spread of a story) and (2) it allows us to examine the effect of an "external forcing" of ideological content. Our assumption is analogous to the inclusion of zealots, which have been studied in various opinion models, including voter models [36] and the Deffuant bounded-confidence model [44]. We let M denote the number of media accounts and n_M denote the number of followers

(i.e., out-edges) per media account. We assume that each media account has the same number of followers, with one exception in Sec. VII (where we impose a media follower-number distribution based on data from real-world media outlets). Once we add media accounts, the total number of nodes in a network is $N + M$.

We represent the ideology of each account in a network at discrete time t by $\mathbf{x}_i^t$ (where $i \in \{1, \ldots, N+M\}$), and we take the ideology space to be continuous and bounded, with $\mathbf{x}_i^t \in [-1, 1]^d$. We allow the ideology space to be d-dimensional, so we can be more nuanced than the typical choice in prior research of using $d = 1$ [45]. At time t, each account spreads content (perhaps with their own spin, as has been studied for memes using Facebook data [46]) that reflects its current ideology $\mathbf{x}_i^t$. When it is clear, we will not always write the superscript t.

B. Content updating rule

We update content synchronously at each discrete time step in the following way. Accounts update their content by averaging the ideologies of observed content that is sufficiently close to their own opinion. We say that an account is *receptive* to such content. Accounts include their own ideology at the previous time step when determining their location in ideological space in a time step. During this step, they share content with their new ideology. Accounts see content only from accounts that they follow; therefore, media accounts are not receptive to any other content. We define a receptiveness parameter c, which gives the distance in ideology within which an account is receptive. With this notation, we define the set of accounts to which account i is receptive to be $I_i = \{j \in \{1, \ldots, N+M\} | A_{ij} = 1; \mathrm{dist}(\mathbf{x}_j, \mathbf{x}_i) < c\}$. We also have to decide how to measure the distance $\mathrm{dist}(\cdot, \cdot)$ in ideological space. One can consider any metric, and we choose to use the ℓ_p norm.

The update rule for the ideology of an account is

$$\mathbf{x}_i^{t+1} = \frac{1}{|I_i| + 1}\left(\mathbf{x}_i^t + \sum_{j \in I_i} \mathbf{x}_j^t\right), \tag{1}$$

which we can also write as

$$\mathbf{x}_i^{t+1} = \frac{1}{|I_i| + 1}\left(\mathbf{x}_i^t + \sum_{j=1}^{N+M} A_{ij}\mathbf{x}_j^t f\left(\mathbf{x}_j^t, \mathbf{x}_i^t\right)\right), \tag{2}$$

where $f(\mathbf{x}) = 1$ if $\mathrm{dist}(\mathbf{x}_j, \mathbf{x}_i) < c$ and $f(\mathbf{x}) = 0$ otherwise. In Fig. 1, we show a schematic of this update rule with a small, concrete example. Our model builds on previous work in opinion dynamics. Our consensus-forming mechanism from averaging ideology is reminiscent of the influential work of DeGroot [47]. Although this is a standard choice, a recent paper by Mei *et al.* [48] suggested that weighted-median influence is also a reasonable choice for consensus formation. In the present paper, we draw our bounded-confidence mechanism from the Hegselmann–Krause (HK) [33] and Deffuant [32,45] models of continuous opinion dynamics. Bounded-confidence updates are also related to the DeGroot–Friedkin (DF) model of social power [49].

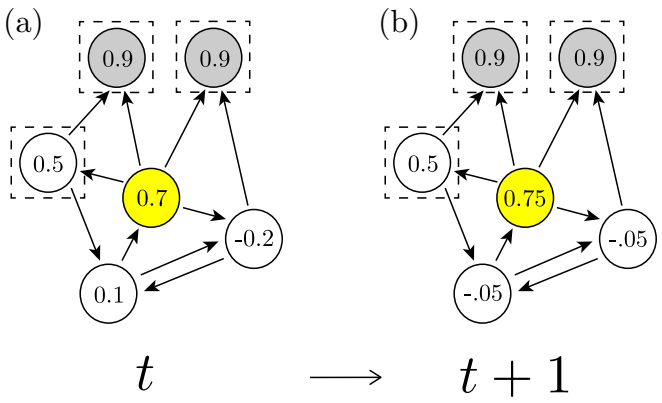

FIG. 1. A schematic of the content updating rule in our model. In (a), we show the ideology of nodes in a social network at time t. In (b), we show their ideologies in the next time step (after we have applied the update rule). In this example, there are $M = 2$ media nodes (gray circles) and $N = 4$ nonmedia nodes. To make this example more concrete, let's focus on the highlighted node (in yellow). The yellow node follows four other accounts, two of which are media accounts and two of which are nonmedia accounts. If we take the receptiveness parameter to be $c = 0.5$, the yellow node is receptive only to content from three of these accounts (specifically, the ones that are surrounded by dashed boxes). After averaging the content from the accounts to which it is receptive, the yellow account updates its ideology for the next time step to $x^{t+1} = \frac{1}{4}(0.7 + 0.5 + 0.9 + 0.9) = 0.75$. During time step $t + 1$, this account shares content with that ideology.

IV. DYNAMICS WITH A ONE-DIMENSIONAL IDEOLOGY

As an example, we examine content spreading when our model has one ideological dimension, where we interpret the value $x_i \in [-1, 1]$ as the political viewpoint of the content of account i on a liberal–conservative axis. We take $x_i = -1$ to be very liberal, $x_i = 1$ to be very conservative, and $x_i = 0$ to be moderate. We use the distance $\mathrm{dist}(x_j, x_i) = |x_j - x_i|$. We then write the content updating rule for these evolving pieces of one-dimensional political content [see Eq. (2)] by writing

$$f(x_j, x_i) = \begin{cases} 1\,, & \text{if } |x_j - x_i| < c \\ 0\,, & \text{otherwise}\,. \end{cases} \tag{3}$$

In this example, we suppose that there are M media accounts that all have the same political opinion $x_M \in [-1, 1]$. There are multiple possible interpretations of this assumption. For example, one can interpret it as M different media outlets that share the same message, which has a given political bias (as represented by a particular location in ideological space). Another way to interpret it is as one media account with $M - 1$ affiliated sockpuppet accounts, which it is using to help spread its message with a specified political bias.

A. Simulations

In each network, we suppose that there are N nonmedia accounts, and we vary the number M of media accounts and the number n_M of followers of each media account. We assume that all media accounts in a given simulation have an equal number of followers, although it would be interesting to extend our model by drawing n_M from a distribution of

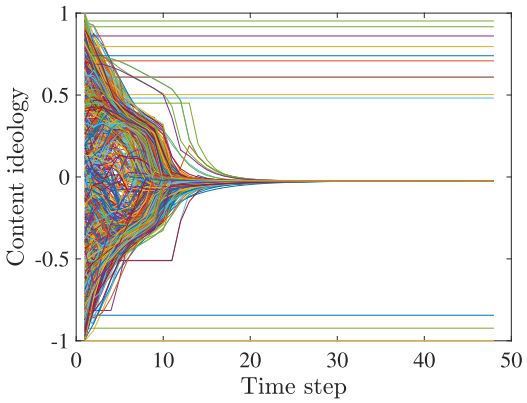

(a)

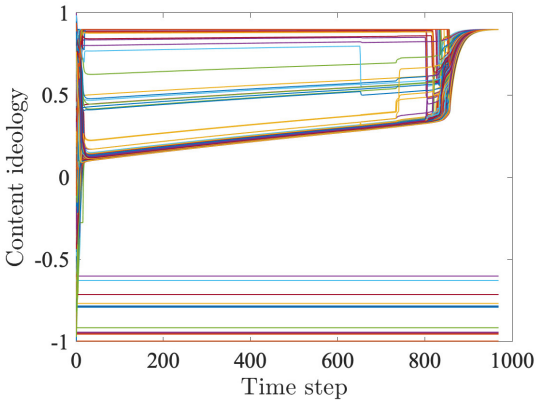

(b)

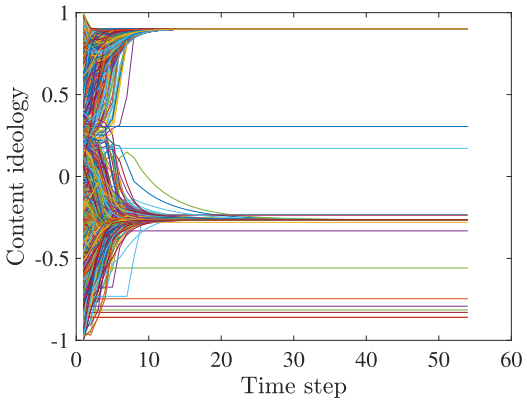

(c)

FIG. 2. Individual trials of our model on the Reed College Facebook network (which has $N = 962$ nonmedia nodes) with different levels of media influence. We influence the convergence time and content-spreading dynamics by varying the media-account parameters M and n_M. (a) With low media influence, most accounts converge to the location $x = 0$ in ideological space. (b) For larger values of M and n_M, a large fraction of the nonmedia accounts converge to the ideological position of the media nodes. (c) For still larger values of these parameters, splitting in ideology occurs very rapidly, resulting in a smaller fraction of nonmedia accounts that converge to the media's ideology. In each simulation in this figure, we evenly distribute the initial ideologies of nonmedia accounts and we set the ideological position of all media accounts to $x_M = 0.9$.

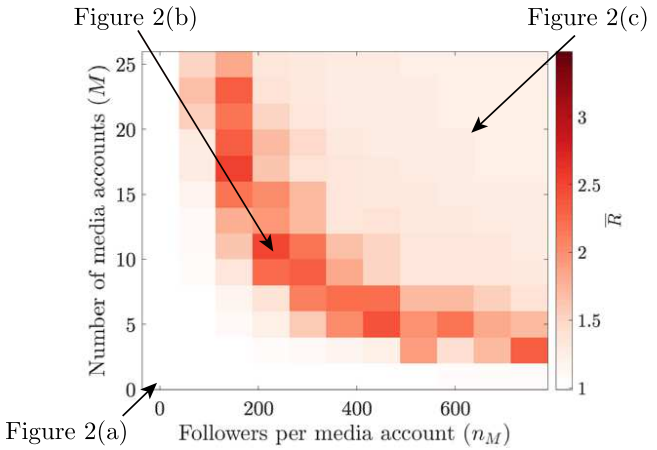

FIG. 3. Media impact for our model on the Reed College network (which has $N = 962$ nonmedia nodes) with different numbers of media accounts (M) on the vertical axis and followers per media account (n_M) on the horizontal axis. All media accounts have the same ideological position of $x_M = 0.9$. The colors indicate the values of $\overline{R}$, the mean of the impact summary diagnostic [see Eq. (7)] for distance from media ideology over 50 trials. Dark red indicates the most media impact (i.e., the largest values of $\overline{R}$), and white indicates the least impact. The ideological position x_M of the media accounts and the initial ideological positions of the nonmedia accounts are the same as those in Fig. 2. The arrows designate the regions of the plot that correspond to the mean media impact for the three panels in Fig. 2.

the numbers of followers. Unless we note otherwise, we use 200 trials for each numerical experiment, and we interpret a trial to have "converged" if $|x_i^t - x_i^{t-1}| < 10^{-4}$ for all i for ten consecutive steps. For simplicity, we set the receptiveness parameter to take the value $c = 0.5$, except for Fig. 8, where we examine the effects of varying this parameter. For a detailed discussion on the effects of the receptiveness parameter in bounded-confidence models, see Meng *et al.* [45].

We initialize the nonmedia accounts to have evenly-spaced initial ideologies, so

$$\mathbf{x}_0 = \left[-1, -1 + \frac{2}{N-1}, \ldots, 1\right].$$

We uniformly randomly permute the starting values of nonmedia nodes for each trial, so they are not spatially ordered in the network. In each trial, each media account has n_M distinct followers, which we choose uniformly at random from the nonmedia accounts. We choose the followers of each media account independently, so it is possible for a nonmedia account to follow multiple media accounts.

In Fig. 2, we show examples of individual trials with different numbers of media accounts and different numbers of followers per media account. We run our model on the Reed College Facebook network from the FACEBOOK100 data set [50]. The networks in this data set are Facebook friendship networks on university campuses from one day in fall 2005. In our simulations, we use only the largest connected component of each Facebook network to represent nonmedia accounts, and we note that the friendship connections between nonmedia accounts in these networks are bidirectional. In

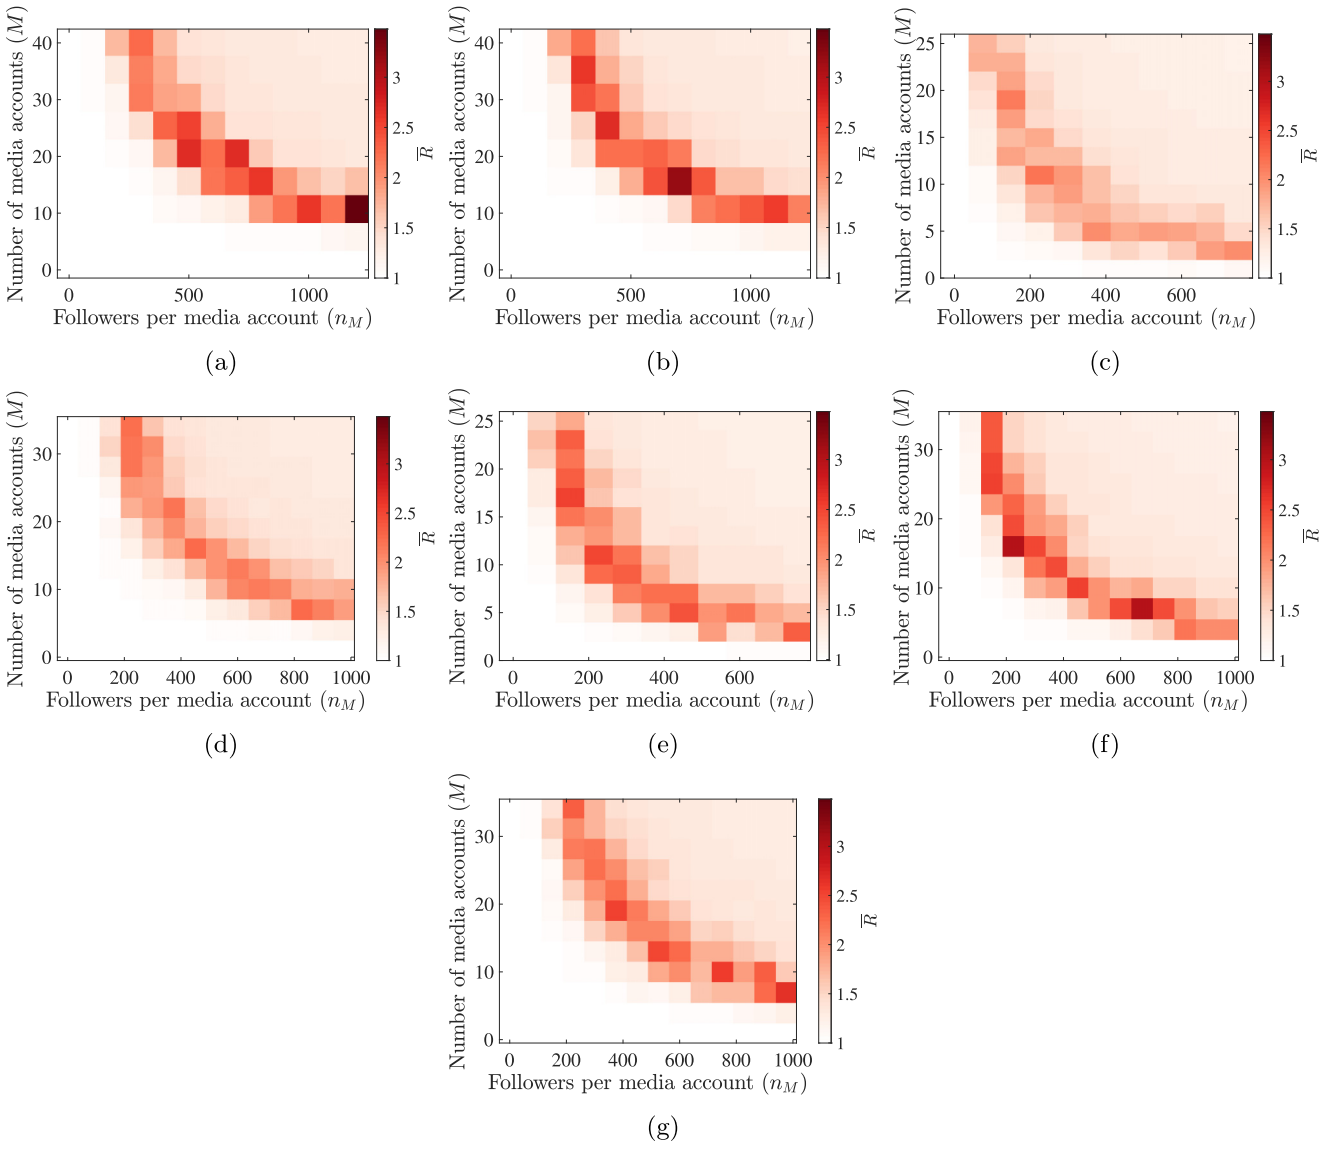

FIG. 4. Heat maps of the summary diagnostic $\overline{R}$ for media impact in the Facebook networks of various universities from the FACEBOOK100 data set [50] illustrate that the largest media impact occurs for moderate numbers of media accounts with moderate numbers of followers. In each panel, the vertical axis gives the number of media accounts (M), the horizontal axis gives the number of followers per media account (n_M), and the colors indicate the mean impact summary diagnostic ($\overline{R}$) over 50 trials. Dark red indicates the most media impact (i.e., the largest values of $\overline{R}$), and white indicates the least impact. The Facebook networks are (a) Amherst, (b) Bowdoin, (c) Caltech, (d) Haverford, (e) Reed, (f) Simmons, and (g) Swarthmore.

Fig. 2, we show the ideology of each of the Reed College network's $N = 962$ nonmedia accounts on the vertical axis and the simulation time t on the horizontal axis. We take the media opinion to be $x_M = 0.9$. These trials suggest that media nodes influence both convergence time and content-spreading dynamics. We explore these ideas further in Sec. IV B.

B. Media influence with a one-dimensional opinion space

Our simulations of our content-spreading model suggest that the number of nonmedia accounts that converge to the ideological position of the media nodes—as we just illustrated for an example (see Fig. 2) in which all media nodes have the same ideology—depends nontrivially both on the number of media accounts and on the number of followers per media account. In this section, we present simulations that illuminate the level of media influence on a variety of real and synthetic networks. We continue to assume that all media nodes have the same ideological position.

We introduce an order parameter that we use as a diagnostic to measure the impact of media nodes on the ideological positions at convergence. Let x_i^b denote the ideological position of account i at convergence in the absence of media influence. We establish a mean baseline ideology R_0 by computing the mean distance between the ideological positions (x_i^b) of nonmedia accounts at convergence and the media ideology x_M. In mathematical terms, this baseline is

$$R_0 = \frac{1}{N}\sum_{i=1}^{N} \left\| x_i^b - x_M \right\|_2 . \tag{4}$$

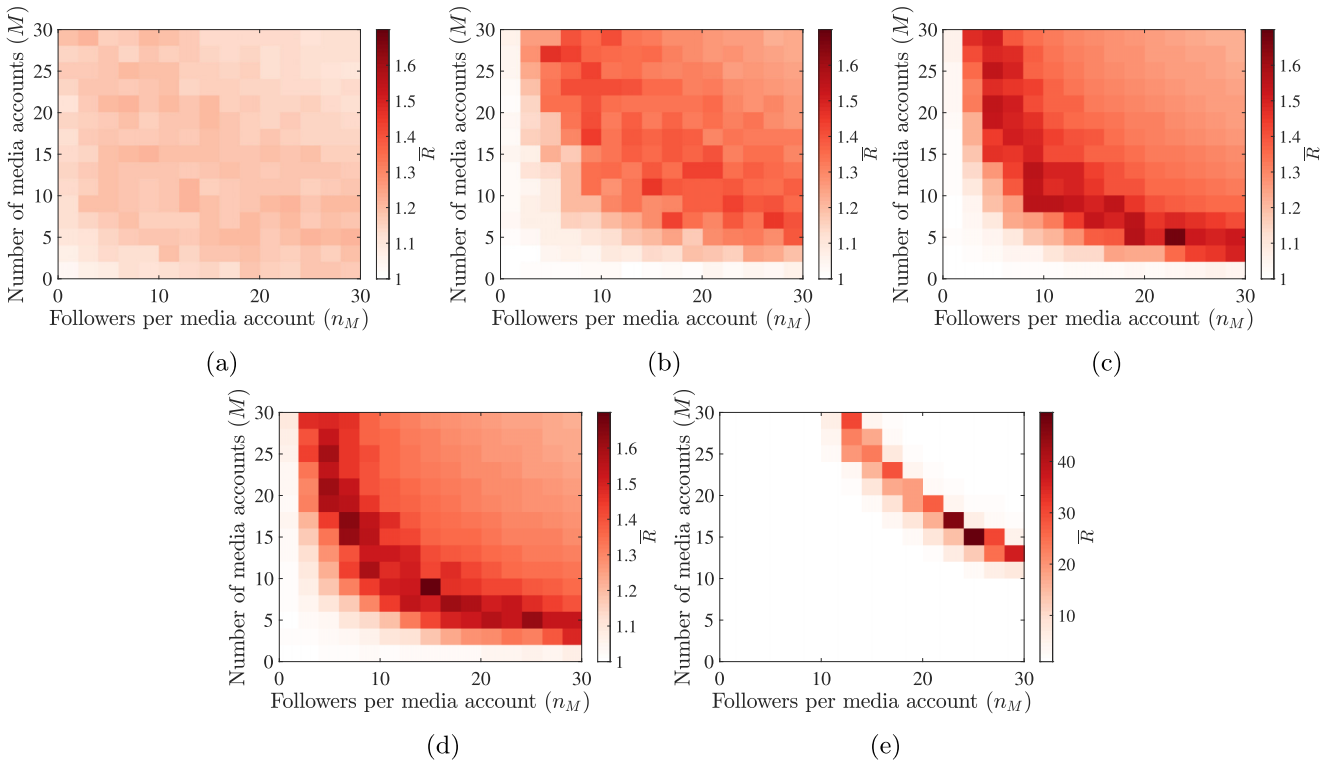

FIG. 5. Heat maps of the summary diagnostic $\overline{R}$ for media impact of our model on various synthetic networks illustrate that the largest media impact occurs for moderate numbers of media accounts with moderate numbers of followers for a variety of network architectures. In each panel, the vertical axis gives the number of media accounts (M), the horizontal axis gives the number of followers per media account (n_M), and the colors indicate the media impact diagnostic ($\overline{R}$), which we compute as a mean over 200 trials. Dark red indicates the most media impact (i.e., the largest values of $\overline{R}$), and white indicates the least impact. In (a), we use a star network in which all nonmedia accounts follow one central nonmedia account, which also follows the peripheral nonmedia accounts. In (b), we use a directed ring lattice in which each nonmedia account follows $k = 25$ neighboring nonmedia accounts. In (c), we construct a directed variant of a Watts–Strogatz network from the directed ring lattice with $k = 25$. In it, we rewire follower connections with a probability of $\beta = 0.5$. In (d), we use a directed variant of an Erdős–Rényi network in which we take the expected mean out-degree (i.e., the mean number of nonmedia accounts that are followed by a nonmedia account) to be $k = 25$. In (e), we use a complete network. For the directed ER and WS networks, we construct a new network for each of the 200 trials.

The quantity R_0 characterizes the mean effect on nonmedia nodes, for one trial, of the dynamics for a given network structure in the absence of media influence. Because the outcome depends on the initial ideological positions of the nonmedia accounts, we average R_0 over many trials.

Once we have calculated the baseline ideology R_0, we construct a similar order parameter R_M that characterizes the mean outcome of the dynamics for the same network after we introduce media nodes. Let x_i^* denote the ideological position of account i at convergence in a network with media influence. The associated media-influenced ideology diagnostic is

$$R_M = \frac{1}{N}\sum_{i=1}^{N} \|x_i^* - x_M\|_2 \,. \tag{5}$$

The order parameters R_0 and R_M allow us to quantify the impact of media nodes on content-spreading dynamics, where we note that one can also define time-dependent analogs of (4) and (5). We define the media impact R for one trial to be the ratio of the mean baseline ideological distance to the media-influenced ideological distance in that trial:

$$R = \frac{\overline{R_0}}{R_M} \,. \tag{6}$$

We can also obtain an overall media impact

$$\overline{R} = \frac{\overline{R_0}}{\overline{R_M}} \tag{7}$$

by averaging the media-influenced opinion function R_M over some number (e.g., 200) of trials. We interpret the media impact in the following way. If $\overline{R} = 1$, the media has not had an influence on the "average" (specifically, the mean) ideological position of the nonmedia accounts (and the content that they spread) at convergence. If $\overline{R} \in [0, 1)$, the media nodes have driven the mean ideological position of the accounts to be farther from the media's ideological position than is the case without the media accounts. Finally, if $\overline{R} > 1$, the mean ideological position of the nonmedia accounts (and hence of the content that they spread) is closer to the media ideology than it would be without the media accounts, with larger values of $\overline{R}$ indicating a stronger impact.

Equation (7) gives one of multiple possible summary diagnostics to measure media influence. Another option is to measure the mean of the distances between the media and nonmedia ideologies; that is, one can use only the expected value of R_M from Eq. (5). This alternative entails

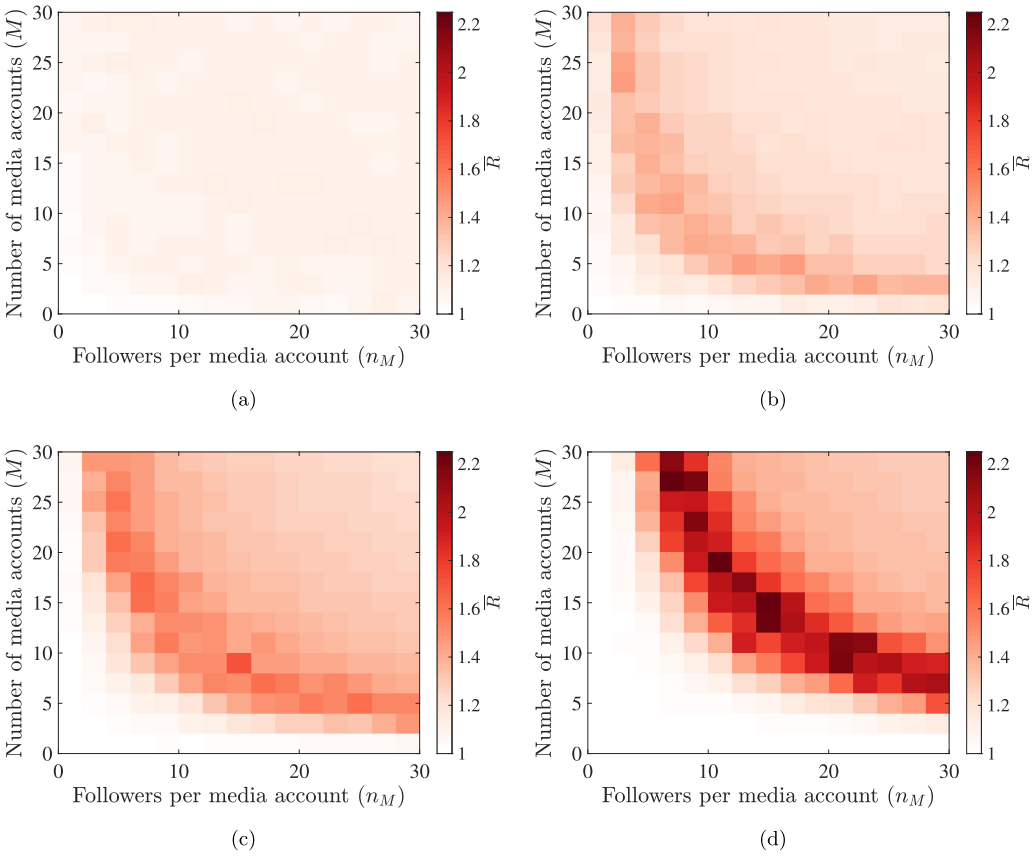

FIG. 6. Heat maps of the media impact diagnostic $\overline{R}$ for directed ER networks with different values of the expected mean number k of nonmedia accounts that nonmedia accounts follow. We show results for (a) $k = 2$, (b) 10, (c) 25, and (d) 50. Our simulations suggest that media impact is larger for progressively larger k. The vertical axis is the number of media accounts (M), the horizontal axis is the number of followers per media account (n_M), and the colors represent the media impact diagnostic ($\overline{R}$), which we average over 200 trials. Dark red indicates the most media impact (i.e., the largest values of $\overline{R}$), and white indicates the least impact.

an impact level of $\overline{R_M} \in [0, 2]$, where small values of $\overline{R_M}$ represent considerable "entrainment" to the media ideology. That is, the media nodes have successfully influenced (e.g., "brainwashed," to describe it in an uncharitable way) many nonmedia accounts in a network. Large values of $\overline{R_M}$ represent low levels of media influence. Using this alternative impact function for the experiments in this section gives the same qualitative results to what we report using (7).

In Fig. 3, we illustrate media impact values for different numbers of media accounts (M) and numbers of media followers (n_M). We show the mean value of the summary diagnostic $\overline{R}$ over 50 trials. These simulations suggest a surprising result: the highest levels of adoption of media ideology do not occur for the largest values in (n_M, M) parameter space. Instead, the most successful scenario for promoting widespread adoption of the media ideology is to spread the content through a moderate number of media accounts, each of which has a moderate number of followers. This observation is consistent with previous empirical observations that accounts with a small number of followers can significantly promote the spreading of content on Twitter [51]. For a very small number of media nodes (or if the media nodes have very few followers), as in Fig. 2(a), the media ideology has very low impact. A large number of media nodes with many followers per account [as in Fig. 2(c)] produces some impact, yielding values of $\overline{R}$ in the interval $(1, 2)$. However, we observe a larger impact when there are a moderate number of media accounts that each have a moderate number of followers [see Fig. 2(b)]. In this situation, we often obtain $\overline{R} \geqslant 2$.

In Fig. 4, we show heat maps of the media entrainment values (using the summary diagnostic $\overline{R}$) for our model for social networks from the FACEBOOK100 data set [50]. Each such network consists of people at a particular university who are connected to each other via Facebook friendships (which yield bidirectional edges) from one day in fall 2005. We add media accounts to each network in the way that we described previously for the Reed College network in Sec. IV A. For each of the FACEBOOK100 networks that we examine, our simulations produce a distinct region with large media impact (specifically, with $\overline{R} \geqslant 2$). In each case, the largest amount of media impact does not occur for the largest values of the two quantities in (n_M, M) parameter space.

In Fig. 5, we show heat maps of the media entrainment values (using the summary diagnostic $\overline{R}$) for our model for a variety of synthetic network architectures, each of which has 100 nonmedia accounts. For clarity, we first briefly describe how we construct each of these networks. These networks are a directed ring lattice, a star, a directed variant of Erdős–Rényi (ER) networks, a directed variant of Watts–Strogatz (WS)

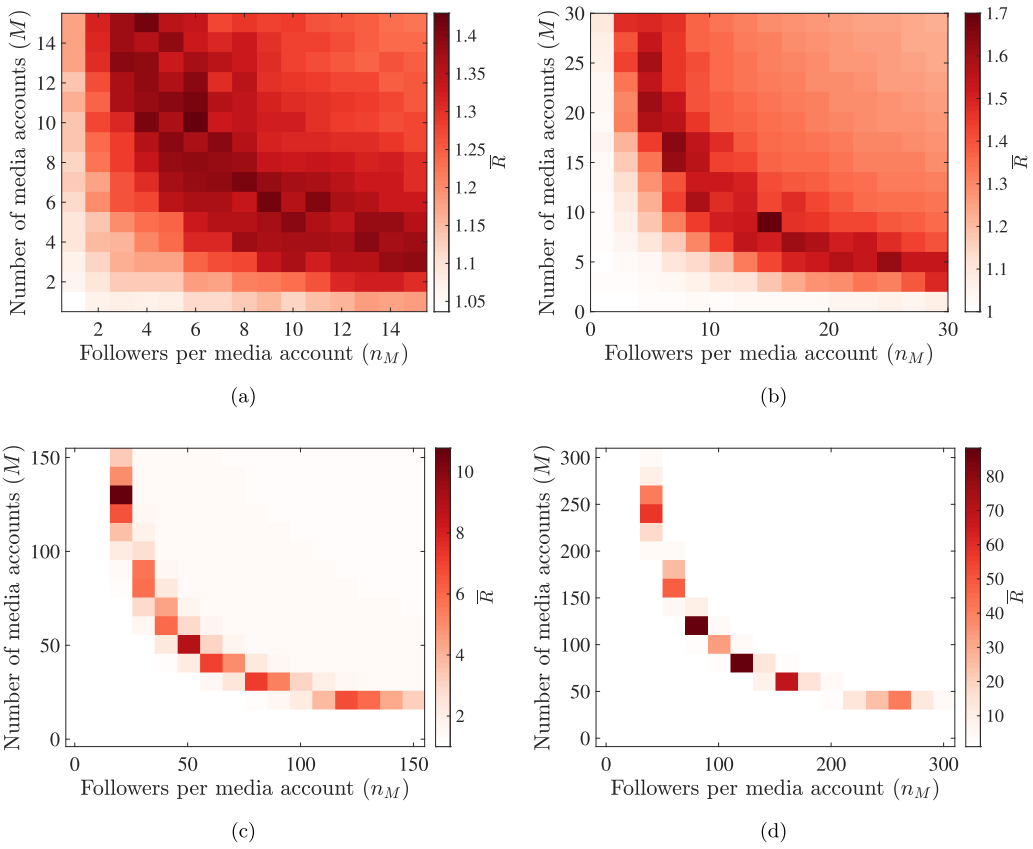

FIG. 7. Heat maps of the media impact diagnostic $\overline{R}$ for directed ER networks with different numbers of nonmedia accounts (N) for a fixed ratio (k/N) of the expected mean number of nonmedia accounts that nonmedia accounts follow to the total number of nonmedia accounts. We show results for (a) $N = 50$ and $k = 13$, (b) $N = 100$ and $k = 25$, (c) $N = 500$ and $k = 125$, and (d) $N = 1000$ and $k = 250$. Our simulations suggest that both the media impact $\overline{R}$ and the spread of the media impact (which we define to be $(\max\{\overline{R}\} - \min\{\overline{R}\})$ over all (n_M, M) pairs) increase with N. The vertical axis is the number of media accounts (M), the horizontal axis is the number of followers per media account (n_M), and the colors represent the media impact diagnostic ($\overline{R}$), which we average over 200 trials. Dark red indicates the most media impact (i.e., the largest values of $\overline{R}$), and white indicates the least impact. Unlike in our other figures, we have not scaled the color range to be the same across these four plots; this aids in visualizing the difference in the spread of impact values for networks of different sizes.

networks, and complete networks. In the star network, all but one of the nonmedia accounts follow one central nonmedia account. These edges are bidirectional, so the central nonmedia account follows all of the other nonmedia accounts. We construct a directed ring lattice as follows. Each nonmedia account follows k other nonmedia accounts, where the ith node follows nodes $\{i+1, i+2, \ldots, i+k\}$ (mod N). In this case, the edges are not bidirectional (so, for example, account 1 follows account 2, but account 2 does not follow 1). We construct a directed variant of a WS network by starting with this directed ring lattice and rewiring each edge with probability $\beta = 0.5$ [52]. We rewire a directed edge from node i to node j by selecting a new node j' uniformly at random from the nodes that account i does not currently follow; we remove the edge from i to j and add an edge from i to j'. We construct directed ER networks, which have an expected mean out-degree of $k = 25$, as a variant of the $\mathcal{G}(N, p)$ model. (In other words, the expected mean number of nonmedia accounts that a node follows is $k = 25$.) For the directed ER and WS networks, we construct a new network for each trial, so one should interpret our simulation results as a sample mean over many networks. Although all of these networks have the same number of nonmedia accounts, the region of parameter space in which the media has the most impact differs across different networks. This suggests that network architecture plays an important role in determining the level of media involvement that is necessary for media accounts to exert "global" influence in a network.

We now examine the media impact on directed ER networks as we vary k, the expected mean number of nonmedia accounts that a nonmedia account follows. In Fig. 6, we show heat maps of the media impact summary diagnostic $\overline{R}$ for directed ER networks with different expected mean out-degrees. As we increase the expected mean out-degree k, we observe that media impact also increases.

We now perform numerical experiments in which we vary the size of the directed ER networks by varying the number of nonmedia accounts. To try to isolate the effect of network size, we fix the ratio (k/N) of the expected mean number (k) of nonmedia accounts that a nonmedia account follows to the total number (N) of nonmedia accounts for each simulation. As in our earlier observations, we see in Fig. 7 that the media impact increases as we increase the number N of nonmedia nodes. However, in this case, we observe a progressively larger spread in the media entrainment diagnostic $\overline{R}$ for progressively larger N. For example, for $N = 1000$ and $k = 250$,

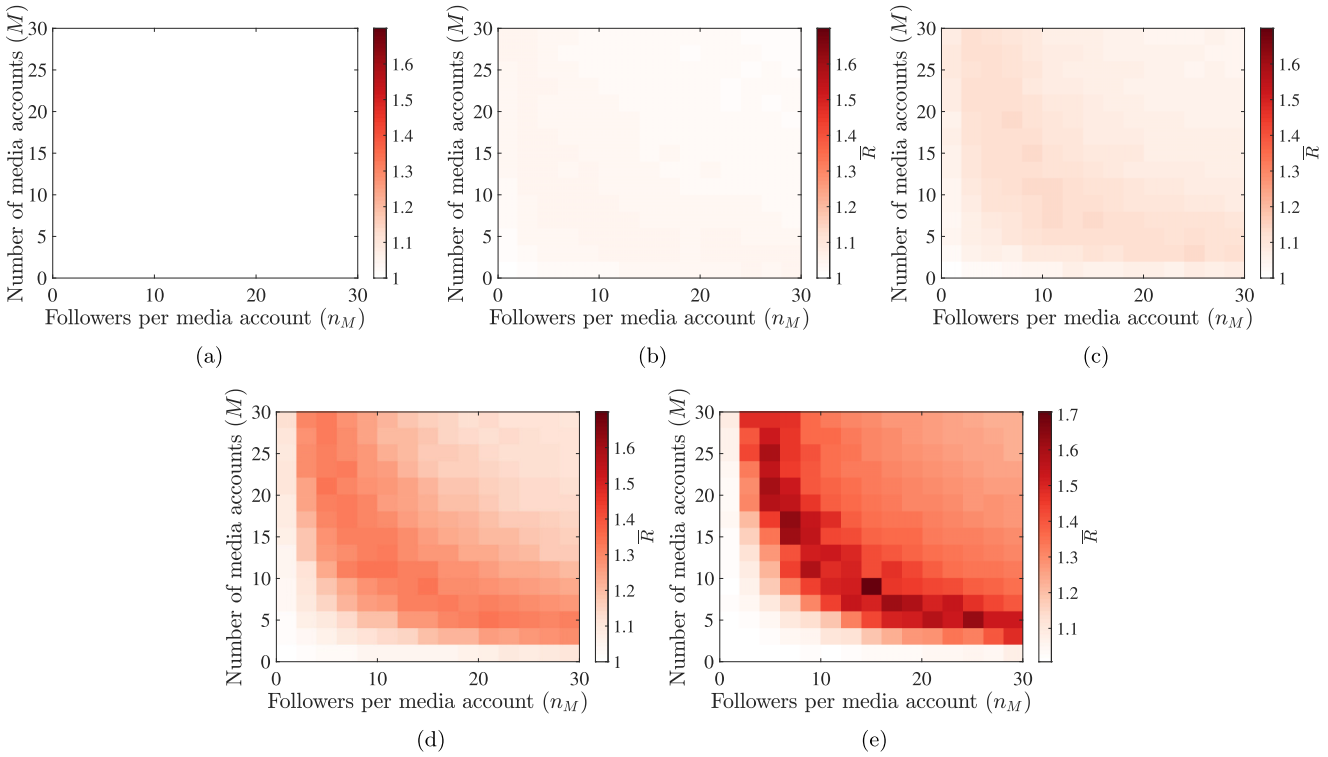

FIG. 8. Heat maps of the media impact diagnostic $\overline{R}$ for directed ER networks for different values of the receptiveness parameter c. We show results for (a) $c = 0.1$, (b) $c = 0.2$, (c) $c = 0.3$, (d) $c = 0.4$, and (e) $c = 0.5$. For small values of c, the receptiveness of accounts is too small for the media accounts to impact the mean ideological positions of the nonmedia accounts. For progressively larger values of c, we observe progressively larger media impact. The vertical axis is the number of media accounts (M), the horizontal axis is the number of followers per media account (n_M), and the colors represent the media impact diagnostic ($\overline{R}$), which we average over 200 trials. Dark red indicates the most media impact (i.e., the largest values of $\overline{R}$), and white indicates the least impact.

the spread of the media impact (which we define to be $(\max\{\overline{R}\} - \min\{\overline{R}\})$ over all (n_M, M) pairs) is approximately $88.2048 - 0.9980 = 87.2068$. However, for $N = 50$ and $k = 13$, the spread is approximately $1.4295 - 1.0364 = 0.3931$ in our simulations.

In Fig. 8, we illustrate the effect of varying the receptiveness parameter c. If the receptiveness is very small, nonmedia accounts in a network adjust their ideological positions only if the content that they see is very close to their current ideology. Consequently, we observe little or no media impact on the ideological positions of the nonmedia nodes in the network. For sufficiently large receptiveness, however, the media do impact the mean ideological position of nonmedia nodes, with a larger impact for larger values of c.

One natural generalization of our model is to consider the effect of varying the weighting of an individual's current

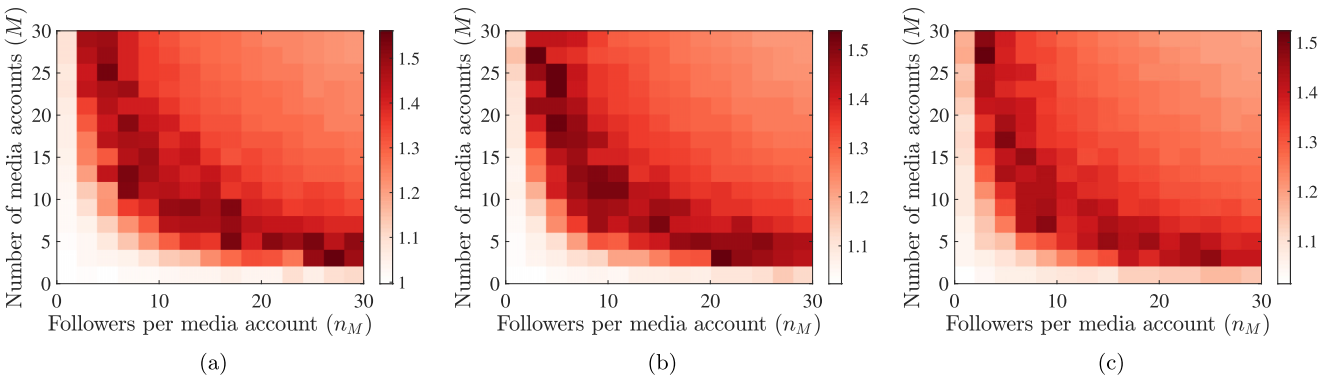

FIG. 9. Heat maps of the media impact diagnostic $\overline{R}$ for directed ER networks for different values of the self-weight: (a) $w = 3$, (b) $w = 5$, and (c) $w = 10$. In our numerical experiments, even when accounts weight their own ideology more heavily than those of the accounts that they follow, we observe similar qualitative dynamics as when we do not incorporate the parameter w into our model. The vertical axis is the number of media accounts (M), the horizontal axis is the number of followers per media account (n_M), and the colors represent the media impact diagnostic ($\overline{R}$), which we average over 200 trials. Dark red indicates the most media impact (i.e., the largest values of $\overline{R}$), and white indicates the least impact.

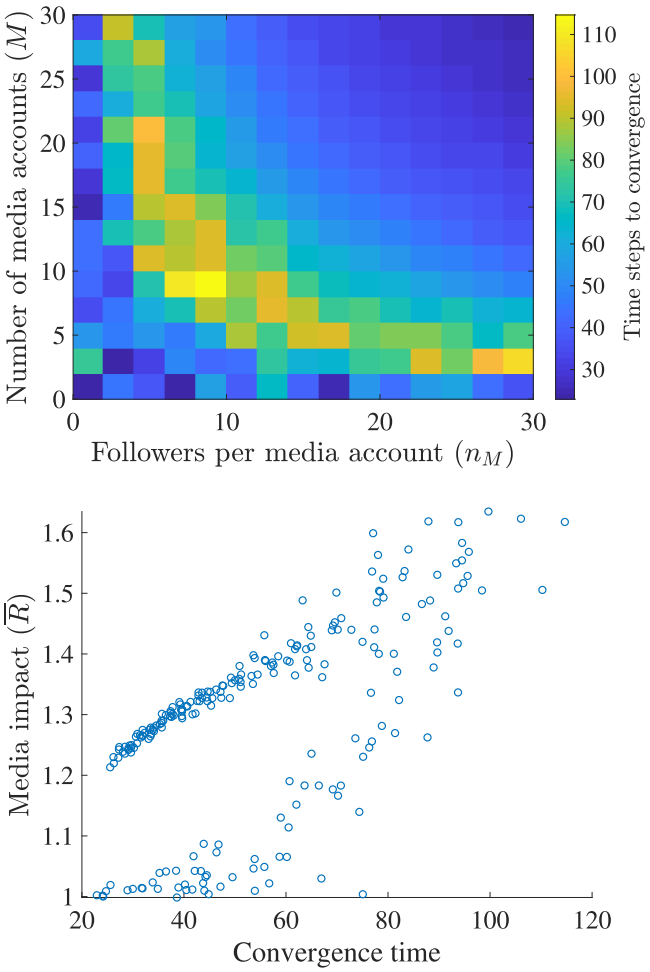

FIG. 10. In our simulations, we observe an interesting relationship between media impact and convergence time. In the top panel, we show a heat map of the mean convergence time (in color, where yellow is the longest time and blue is the shortest) versus the number of media accounts (M) on the vertical axis and the number of followers per media account (n_M) on the horizontal axis. In the bottom panel, we show the mean media impact $\overline{R}$ versus the mean convergence time. Each circle represents one value in the (n_M, M) parameter plane, so each circle corresponds to one of the square regions from (a). A larger media impact is positively correlated with a longer time to convergence. In both panels, we run our simulations on a directed ER network with $N = 100$ nonmedia accounts that follow an expected mean of $k = 25$ nonmedia accounts.

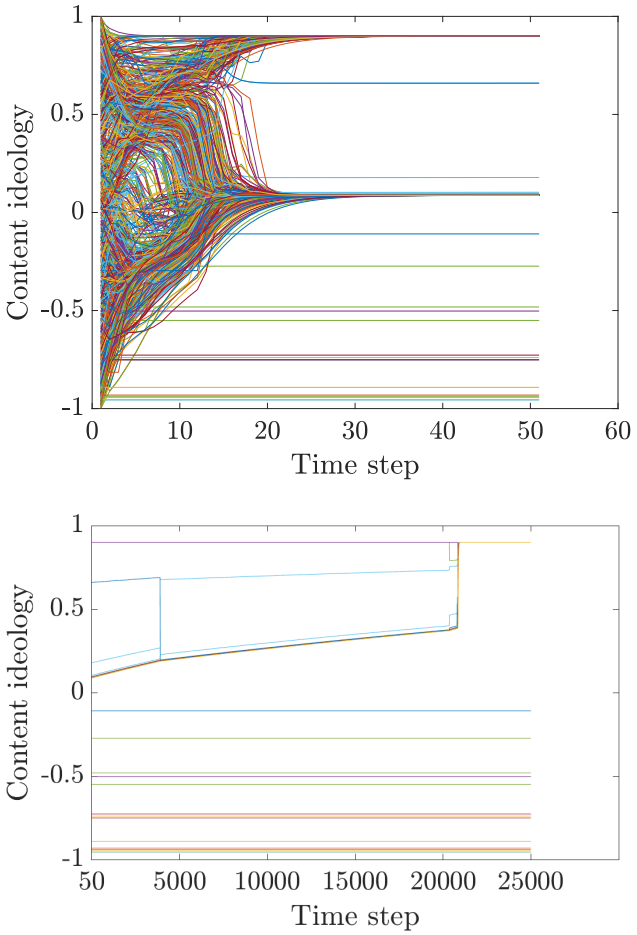

FIG. 11. An example in which our model exhibits multiple timescales in a single simulation on the Reed College Facebook network (which has $N = 962$ nonmedia nodes) with the same parameter values ($c = 0.5$ and $x_M = 0.9$) as in Fig. 3. In this simulation, $M = 9$ and $n_M = 225$. In the top panel, we set the convergence tolerance to TOL $= 10^{-4}$; our simulation converges with this tolerance after 51 time steps. In the bottom panel, we show the same network with the same initial data with the convergence tolerance set to TOL $= 0$. For this convergence tolerance, the two large ideology groups with positive ideologies eventually "collapse" on a long timescale to one group with a single ideology, which is entrained to the media's ideology. This "collapse" takes over 2×10^4 time steps to occur; this is several orders of magnitude longer than the dynamics in the top panel.

ideology in its content updating rule. (This is reminiscent of self-appraisal in DeGroot models [31,49].) One way to incorporate such a weighting, with which an account can value its current ideology more than those of the accounts that it follows, is to include a parameter $w \in \mathbb{R}_{\geqslant 0}$ into the update rule in Eq. (2). The content updating rule with self-weight w is

$$\mathbf{x}_i^{t+1} = \frac{1}{|I_i| + w}\left(w\mathbf{x}_i^t + \sum_{j=1}^{N+M} A_{ij}\mathbf{x}_j^t f\left(\mathbf{x}_j^t, \mathbf{x}_i^t\right)\right). \tag{8}$$

We perform numerical experiments on directed ER networks in which we increase the self-weight w such that nonmedia nodes weight their current ideology 3, 5, and 10 times more than the ideologies of the accounts that they follow. We show the results of these simulations in Fig. 9. In these examples, we observe media impact that is qualitatively similar to our observations without the parameter w.

Additionally, as suggested by the individual trials in Fig. 2, we observe a relationship between media impact and time to convergence. As we illustrate in Fig. 10, our simulations indicate that a larger media impact is positively correlated with a longer time to convergence.

C. Metastability and long-time dynamics

In Sec. IV B, we defined a numerical simulation to have "converged" if $|x_i^t - x_i^{t-1}| < \text{TOL} = 10^{-4}$ for all i. Given

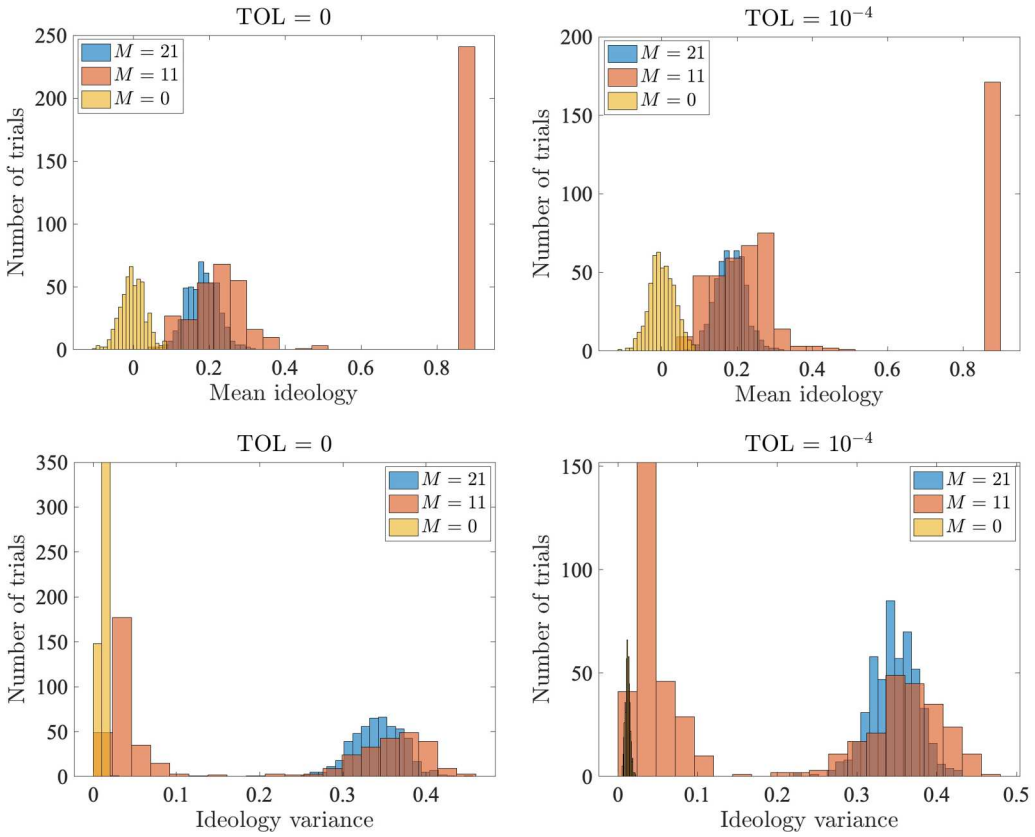

FIG. 12. We simulate our model on the Reed College Facebook network (which has $N = 962$ nonmedia nodes) with the same parameter values ($c = 0.5$ and $x_M = 0.9$) as in Fig. 3 for 500 trials for three levels of media impact: $M = 0$ and $n_M = 0$ (yellow), $M = 11$ and $n_M = 225$ (orange), and $M = 21$ and $n_M = 675$ (blue). In the top row, we show the distributions of the mean ideology at convergence for convergence tolerances of (left) 0 and (right) 10^{-4}. In the bottom row, we show the distributions of the variance of the ideology at convergence for convergence tolerances of (left) 0 and (right) 10^{-4}. Our histograms have 20 bins. (In some cases, the bins are very thin.) We observe that the qualitative behavior of the mean and variance of the ideologies is the same for different convergence tolerances.

this definition, an important question to ask is whether (1) these "converged" values have truly reached a stationary state or (2) our model will subsequently reach a consensus state on extremely long timescales. To examine this question, we run numerical simulations of our model with a convergence tolerance of TOL $= 0$ and impose a bailout time of $T =$ 250 000 time steps. If our simulations reach the bailout time, we record the "converged" state as the state at the bailout time. We perform these simulations on the Reed College Facebook network (which has $N = 962$ nonmedia nodes) with the same parameter values ($c = 0.5$ and $x_M = 0.9$) as in Fig. 3 and on directed ER networks with the same parameter values ($N = 100$, $k = 25$, $c = 0.5$, and $x_M = 0.9$) as in Fig. 5.

Although the apparent "convergence" of our simulations does exhibit interesting dynamics at multiple timescales in some trials (see Fig. 11), these cases are rare. The distributions for the mean have the same form regardless of whether we choose TOL $= 10^{-4}$ or TOL $= 0$. The same is true for the distributions of the variance. In Figs. 12 and 13, we show histograms of the mean ideology and variance of the ideology over 500 trials in the Reed College Facebook network (see Fig. 12) and in directed ER networks (see Fig. 13). Both figures illustrate that the qualitative dynamics are the same for both tolerance levels. In yellow, we show the distributions without media impact ($M = 0$ and $n_M = 0$). In orange, we show the distributions when there are a moderate number of media accounts that each have a moderate number of followers. In blue, we show the distributions when there are a large number of media accounts that each have a large number of followers. As in our simulations in Sec. IV B, the mean ideologies have larger values in the presence of media nodes, and perfect entrainment (for which the mean ideology is 0.9) occurs primarily when there are a moderate number of media accounts (e.g., $M = 15$ for the ER networks), rather than when there are a larger number of media accounts (e.g., $M = 30$ for the ER networks). We supplement these observations by examining the distributions of the variance in ideology. Small variance in ideology corresponds to consensus states, and we observe these states only when media influence is absent (i.e., $M = 0$) or low or moderate (e.g., $M = 15$ for the ER networks). In other words, high levels of media influence (e.g., $M = 30$ for the ER networks) do not result in ideological consensus (including as a result of entrainment to the media ideology), even at 0 tolerance.

In the present paper, we focus on dynamics on short and medium timescales, as we believe that it is particularly relevant to simulate the dynamics of content spread over the course of a news cycle. However, the observed separation of timescales in some trials of our numerical experiments may have interesting implications for future studies on long-term effects of media impact in online social networks.

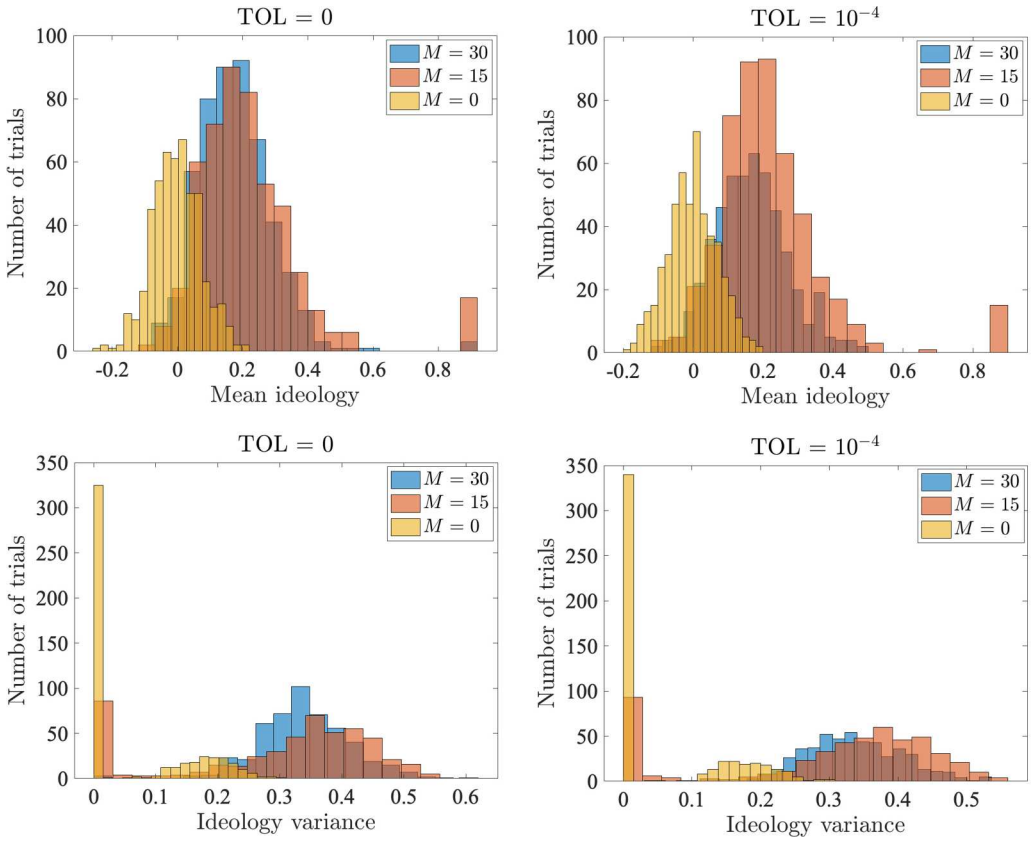

FIG. 13. We simulate our model for 500 trials, where we draw the network for each trial from a directed ER random-graph ensemble with $N = 100$ nodes and an expected mean out-degree of $k = 25$ (as in Fig. 5). We take $c = 0.5$ and $x_M = 0.9$, and we examine three levels of media impact: $M = 0$ and $n_M = 0$ (yellow), $M = 15$ and $n_M = 15$ (orange), and $M = 30$ and $n_M = 30$ (blue). In the top row, we show the distributions of the mean ideology at convergence for convergence tolerances of 0 (left) and 10^{-4} (right). In the bottom row, we show the distributions of the variance of the ideology at convergence for convergence tolerances of 0 (left) and 10^{-4} (right). Our histograms have 20 bins. For both tolerances, we observe that the mean ideology has the same qualitative behavior; the same is true of the variance of the ideology.

V. EXTENDING OUR MODEL TO TWO IDEOLOGICAL DIMENSIONS

We now move beyond one ideological dimension and examine our content-spreading model with two ideological dimensions. Suppose that, at time t, each account has an ideological position of $\mathbf{x}_i \in [-1, 1] \times [-1, 1]$. We use the ℓ_2 norm to calculate the distance $\text{dist}(\mathbf{x}_j, \mathbf{x}_i) = \|\mathbf{x}_j - \mathbf{x}_i\|_2$ in ideology between accounts i and j. Depending on context, this two-dimensional ideology may represent two different aspects of a political-bias spectrum (e.g., the first dimension may represent socially liberal versus socially conservative, whereas the second may represent preferred economic policies that very between socialism and capitalism), or it may represent political bias on multiple issues (e.g., immigration reform and gun control). Given our ideological space and choice of metric on this space, we then write our model as in Eq. (2).

Simulations

We simulate our content-spreading model with two ideological dimensions using a straightforward generalization of the simulations of our model with one ideological dimension. We add media accounts with a given ideology $\mathbf{x}_M$ and assign their followers uniformly at random from the nonmedia accounts. In our examples, we initialize nonmedia accounts with ideological positions that we draw uniformly at random for each trial. (Therefore, it is again true that there is no spatial ordering of ideology in the network structure.)

In Fig. 14, we show a heat map of simulations of media impact in our model with two ideological dimensions. We observe qualitatively similar results as in our simulations of our content-spreading model with one-dimensional ideologies. We again see that the largest media impact (i.e., entrainment) does not occur for the largest numbers of media accounts and followers per media account. Instead, the largest amount of entrainment occurs when there are a moderate number of media accounts, each of which has a moderate number of followers.

VI. COMBINING MEDIA BIAS AND QUALITY

Thus far, we have used our content-spreading model to study the effects of media ideology on content spreading in social networks. We now introduce the primary novel component of our work: incorporation of media content quality into our spreading model. Bessi *et al.* [5] observed that individuals are more likely to spread low-quality content if it confirms or supports their existing biases, and we seek to incorporate such behavior into our content-spreading model.

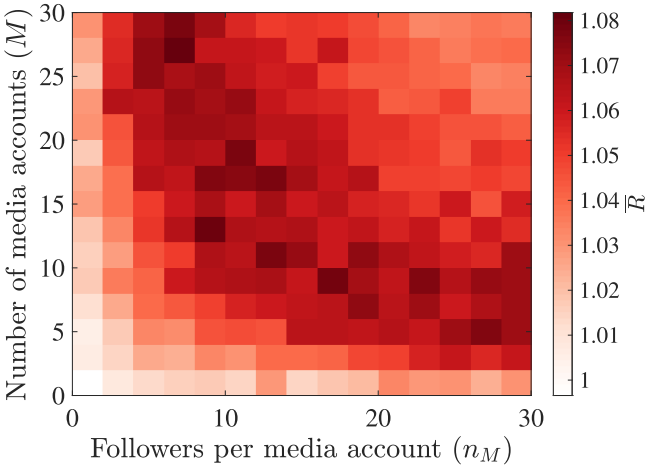

FIG. 14. An example of a heat map of the media impact diagnostic $\overline{R}$ for our model with political positions with two ideological dimensions. For this example, we simulate our model on directed ER networks with $N = 100$ nonmedia accounts that follow an expected mean of $k = 25$ nonmedia accounts. We obtain qualitatively similar results as in our computations with one ideological dimension, but the mean media impact is less pronounced in this case. Specifically, the mean impact $\overline{R}$ is smaller when using two ideological dimensions than what we observed previously when using one ideological dimension for directed ER networks with the same parameter values. The vertical axis is the number of media accounts (M), the horizontal axis is the number of followers per media account (n_M), and the colors represent the media impact diagnostic ($\overline{R}$), which we average over 200 trials. Dark red indicates the most media impact (i.e., the largest values of $\overline{R}$), and white indicates the least impact.

We again denote the state of account i by $\mathbf{x}_i$. In our prior discussions, this state was a vector in d-dimensional ideological space (with $d = 1$ and $d = 2$ in our simulations). Now, however, we introduce an additional dimension into an account's current state to indicate the quality of its shared content; therefore, we now take $\mathbf{x}_i \in [-1, 1]^d \times [0, 1]$. We allow content quality to lie on a continuous spectrum, with values near $x_{d+1,i} = 0$ representing propaganda and disinformation (e.g., clickbait or so-called "fake news") and values closer to $x_{d+1,i} = 1$ representing substantive, thorough, fact-based material.

To focus our discussion, suppose that there is one ideological dimension, so $d = 1$ and $\mathbf{x}_i = (x_{1,i}, x_{2,i}) \in [-1, 1] \times [0, 1]$. As before, suppose that accounts adjust their ideological views only when they are exposed to content that is within a distance c of their current ideology (i.e., only when $\text{dist}(x_{1,i}, x_{1,j}) < c$). Additionally, we now suppose that accounts also consider content quality when we determine their receptiveness. Account i decides whether the content of account j is acceptable for spreading based on the distance $\text{dist}(x_{1,i}, x_{1,j})$ between their ideological positions. If this distance is very small (e.g., close to 0), this content supports account i's ideology, and i is more likely to spread it even if it is not of high quality. However, if the distance in ideology is larger (e.g., close to c), account i is more discerning and tends to spread the content only when it is of sufficiently high quality.

We quantify content discernment in the following way. We calculate the minimum acceptable quality $q_{i,j}$ as a linear function of distance in ideology between account i and account j. That is,

$$q_{i,j} = \frac{1}{c}\text{dist}(x_{1,i}, x_{1,j})\,. \tag{9}$$

With this functional form of $q_{i,j}$, account i spreads content that confirms its bias exactly regardless of quality, because the minimum acceptable quality is 0. If the content has an ideological position that equals the maximum receptiveness distance c from account i's ideological position, then account i spreads the content only if it is of quality 1, the highest possible quality.

The ideological-position updating rule with quality discernment is

$$x_{1,i}^{t+1} = \frac{1}{|I_i| + 1}\left(x_{1,i}^t + \sum_{j=1}^{N+M} A_{ij} x_{1,j}^t g\big(\mathbf{x}_i^t, \mathbf{x}_j^t\big)\right), \tag{10}$$

where

$$g\big(\mathbf{x}_i^t, \mathbf{x}_j^t\big) = \begin{cases} 1\,, & \text{if } x_{2,j} > q_{i,j} \\ 0\,, & \text{otherwise} \end{cases} \tag{11}$$

and

$$I_i = \big\{j \in \{1, \ldots, N+M\} \big| A_{ij} = 1; g(\mathbf{x}_i, \mathbf{x}_j) = 1\big\}. \tag{12}$$

We take the metric to be the ℓ_2 norm $\text{dist}(x_i, x_j) = \|x_{1,i} - x_{1,j}\|_2$, so we are using the same ideology metric as in our previous discussions. When a nonmedia node elects to spread content, we also adjust the quality of the content that it spreads with the update

$$x_{2,i}^{t+1} = \frac{1}{|I_i| + 1}\left(x_{2,i}^t + \sum_{j=1}^{N+M} A_{ij} x_{2,j}^t g\big(\mathbf{x}_i^t, \mathbf{x}_j^t\big)\right), \tag{13}$$

which, along with Eq. (10), implies that

$$\mathbf{x}_i^{t+1} = \frac{1}{|I_i| + 1}\left(\mathbf{x}_i^t + \sum_{j=1}^{N+M} A_{ij} \mathbf{x}_j^t g\big(\mathbf{x}_i^t, \mathbf{x}_j^t\big)\right). \tag{14}$$

VII. MEASURING IMPACT FROM MULTIPLE SOURCES: MEDIA BIAS AND QUALITY

Earlier in our paper (see Sec. IV B), we studied the effect of the numbers of media accounts and media followers per media account when media content has one fixed ideology. Our incorporation of content quality allows us to examine an important question of societal interest in a simple but plausible way: How do competing media accounts influence the outcome of opinions or ideological positions in a social network? This situation models a more realistic scenario of the influence of disparate news media (with heterogeneous political biases and content quality) in a social network. We extend our media-entrainment summary diagnostic from our model with one ideological dimension to measure the impact of media accounts on the mean ideological position at convergence when there are multiple media sources. In this scenario, the quantity of interest is a function $\overline{R}(\mathbf{x})$ that encodes the impact of each possible ideological position $\mathbf{x}$.

To construct the function $\overline{R}(\mathbf{x})$, we first establish a baseline function $r_0(\mathbf{x})$ that encodes the prevalence of each ideology

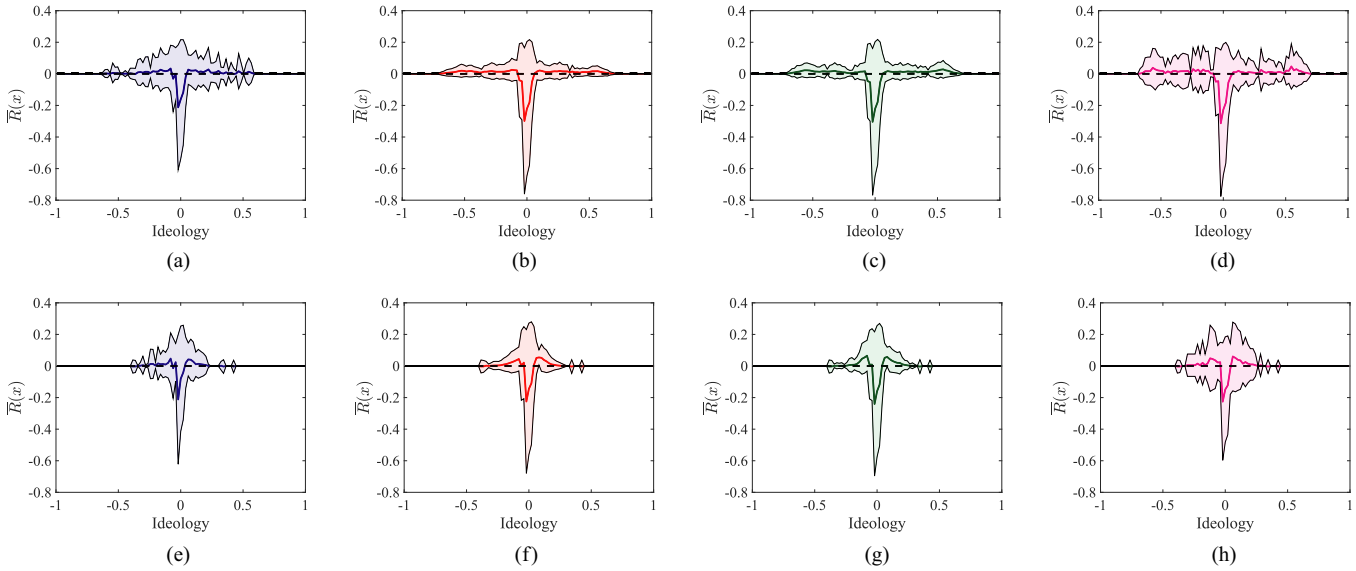

FIG. 15. Impact on the ideology x (which we take to be a scalar in this example) of nonmedia accounts in a network with media ideologies that we draw from a probability distribution. In the top row [i.e., (a)–(d)], we show the media impact for $M = 100$ media accounts, where we draw the ideological positions of each media account from the uniform distribution on $[-1, 1]$. We consider the following numbers of followers per media account: (a) $n_M = 1$ (blue), (b) $n_M = 25$ (red), (c) $n_M = 50$ (green), and (d) $n_M = 100$ (pink). In the bottom row [i.e., (e)–(h)], we show the media impact for $M = 100$ media accounts with ideological positions that we draw from a truncated Gaussian distribution on $[-1, 1]$ with (before truncation) mean 0 and standard deviation 0.5. We again consider the following numbers of followers per media account: (e) $n_M = 1$ (blue), (f) $n_M = 25$ (red), (g) $n_M = 50$ (green), and (h) $n_M = 100$ (pink). Both sets of simulations are on directed ER networks with $N = 100$ nonmedia accounts. We average each media impact function $\overline{R}(x)$ over 200 trials. We indicate the standard deviations with the shaded regions. The dashed line at $\overline{R}(x) = 0$ represents no media impact. Values of $\overline{R}(x)$ that lie below this line indicate that the media has decreased the prevalence of nonmedia accounts with ideology x in comparison to what occurs in the absence of media; values of $\overline{R}$ above this line indicate that the media accounts have increased the prevalence of nonmedia accounts with ideology x.

for a given network in the absence of media accounts. Ideally, it may be desirable to take $r_0(\mathbf{x})$ to be the probability that an account prefers content with ideological position $\mathbf{x}$. In practice, we construct a function $r_0(\mathbf{x})$ by binning the ideology of the content for each trial into bins of width $\delta\mathbf{x}$. We then count the number of times that there is content in the ideology interval $[\mathbf{x}, \mathbf{x} + \delta\mathbf{x}]$, and we calculate a normalized histogram. Once we have constructed the baseline function $r_0(\mathbf{x})$, we use our previous strategy (see Sec. IV B) to construct a function $r_i(\mathbf{x})$ that measures the prevalence of each ideology in a network in the presence of media accounts. That is, $r_i(\mathbf{x})$ is the distribution that describes the probability that a nonmedia account has ideological position $\mathbf{x}$ when there are media accounts in the network. We then use these functions to construct the media impact function $\overline{R}(\mathbf{x}) = r_i(\mathbf{x}) - r_b(\mathbf{x})$. Positive values of $\overline{R}(\mathbf{x})$ for ideology $\mathbf{x}$ indicate that the media has enhanced the prevalence of content with this ideology in a network, negative values indicate that the media has decreased the prevalence of content with this ideology in a network, and 0 indicates no change.

In our model, it is interesting to consider a variety of distributions of media ideologies. For example, we can draw these opinions from a convenient synthetic probability distribution, or we can determine them from empirical data. First, using one ideological dimension, we consider two examples in which we draw media ideologies from synthetic probability distributions. In Fig. 15, we show the media impact functions when there are $M = 100$ media accounts with ideologies that

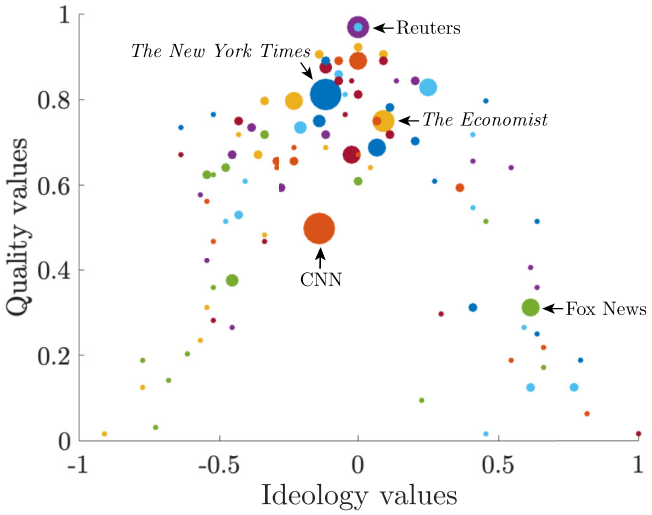

FIG. 16. Distribution of media accounts in (ideology, quality) space. We include $M = 103$ media accounts, whose positions in the space are hand-curated and available from the Ad Fontes Media Bias Chart [53]. Each colored dot represents one media account. The horizontal coordinate represents the ideological position of an account's political content, and the vertical coordinate represents the quality of its political content. The size of a dot represents the number of followers of the associated Twitter account. The five labeled accounts are the ones that are followed by at least 25% of the nonmedia accounts in our numerical experiments.

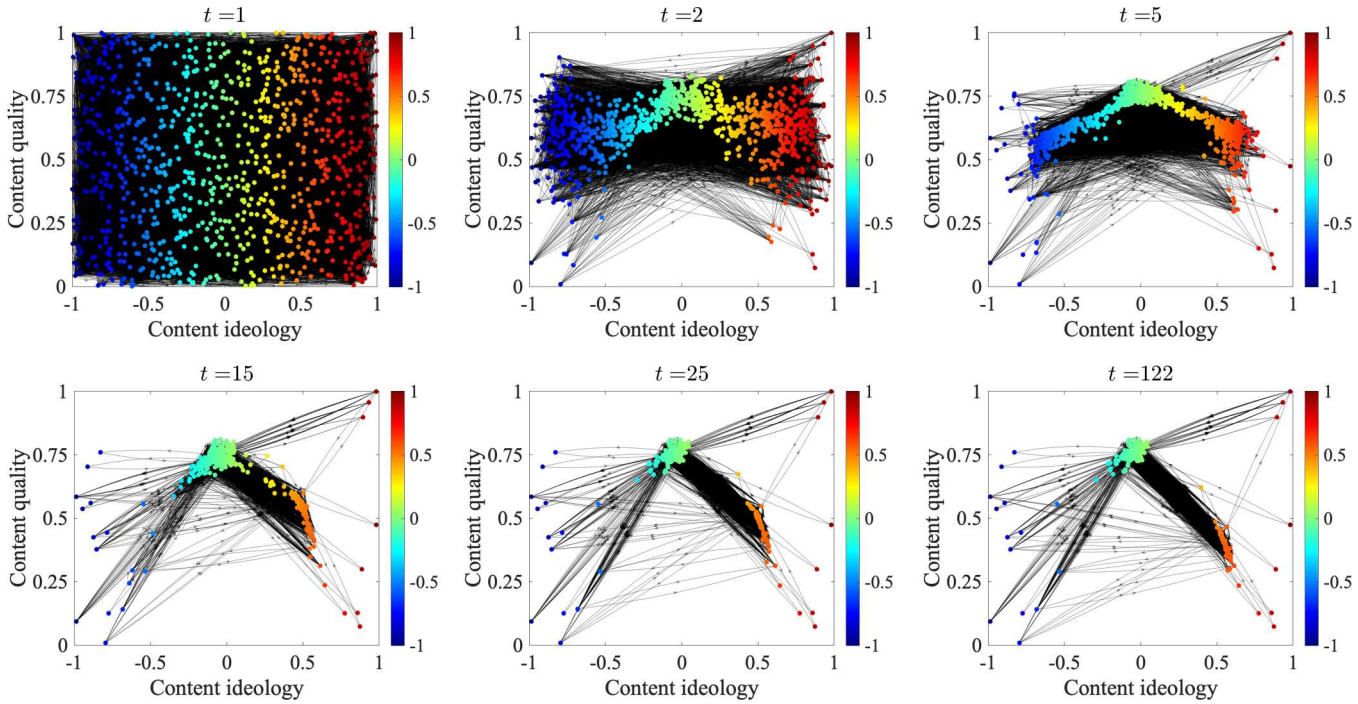

FIG. 17. Snapshots of temporal evolution of content ideology and quality in the Reed College network with $M = 103$ media accounts over 122 time steps. We draw media ideologies and qualities from Version 4.0 of the Ad Fontes Media Bias Chart [53]. We choose the number of followers per media account to be proportional to the approximate number of followers that each media source had on Twitter on 15 February 2019 at 17:36 Pacific Standard Time (under the constraint that each media source in the model has at least one follower, so we round up to one follower for some accounts). Each dot in the figure represents a nonmedia account. The horizontal coordinate represents the ideological bias of each account's content at time t, and the vertical coordinate represents its content quality at time t. To aid visualization, the colors also signify ideological position. (The bluest color is the most liberal, and the reddest color is the most conservative.) We represent follower relationships as lines between the dots. This network architecture does not change over time, although the content that is spread by the nonmedia accounts does change over time. Because the ideology and quality of the content from the media accounts do not change, we do not show these nodes.

we draw (1) from a uniform distribution on $[-1, 1]$ and (2) from a truncated Gaussian distribution on $[-1, 1]$ with (before truncation) mean 0 and standard deviation 0.5. For each of these two examples, we examine the dependence of the media impact function on the number n_M of followers per media account.

We also showcase our model using a media distribution that we generate from hand-curated empirical observations of real-world media sources that incorporate both political ideology (in one dimension) and content quality (in the other dimension). This example includes $M = 103$ media accounts with ideological biases and qualities from Version 4.0 of the hand-curated Ad Fontes Media Bias Chart [53]. We rescale the ideology and quality chart coordinates so that they lie within the intervals $[-1, 1]$ and $[0, 1]$, respectively. In Fig. 16, we show these media content coordinates. These coordinates provide an illustrative example of a possible input to our model; one should not necessarily construe them as quantitatively representing the "true" ideologies of the depicted sources. An important extension of our work is developing quantitative techniques (e.g., using sentiment analysis) to analyze bias and quality of content sources from real data. This is a difficult and interesting problem, and we leave it as future work.

In Fig. 17, we show the temporal evolution of one trial of our content-quality spreading model on the Reed College network with media account ideologies and qualities from the Media Bias Chart. We set each media account j to have a number $n_{M,j}$ of followers that is proportional to its number of followers on Twitter on 15 February 2019 at 17:36 Pacific Standard Time (under the constraint that each media source in the model has at least one follower, so we round up to one follower for some accounts). We select followers for each media account by selecting the nonmedia accounts whose initial conditions are closest in ideological position to that of the media account. That is, for each media account j (with $j \in \{1, \ldots, M\}$) and each nonmedia account i (with $i \in \{1, \ldots, N\}$), we calculate the distance $\mathrm{dist}(x^0_{1,i}, x_{1,M_j})$ in ideological position for all i at $t = 0$. From this set of all distances in ideological position from media account j, we select the n_{M_j} smallest distances; these nonmedia accounts are the followers of media account j. In the simulations for which we use this media distribution as an input, we observe the emergence of two primary communities ("echo chambers") of content: one in which the content is ideologically moderate and of fairly high quality (specifically, it has an ideology of about 0 and a quality of about 0.75) and one in which the opinion is more conservative but of lower quality (specifically, it has an ideology of about 0.5 and a quality of about 0.4). See the contour plot in Fig. 18. The polarization that results from simulations of our model is a reflection of the polarization in the ideology and quality of the media distribution. In our simulations, the account for Fox News (which has an

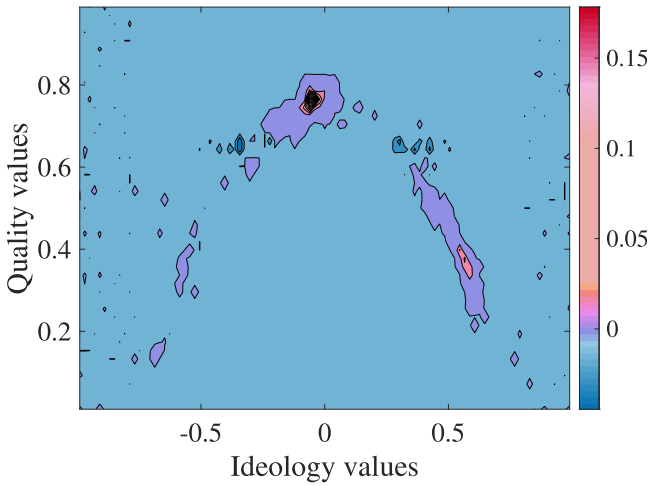

FIG. 18. Contour plot of the media impact function $\overline{R}(\mathbf{x})$ for ideology and quality in the Reed College Facebook network with $M = 103$ media accounts, where we draw media biases and qualities from Version 4.0 of the Ad Fontes Media Bias Chart [53] (as in Fig. 17). The impact function $\overline{R}(\mathbf{x})$ illustrates the relative influence of the media accounts on the mean ideology and quality of content in the network. When $\overline{R} < 0$ (in blue), the media has decreased the prevalence of content with the (ideology, quality) pair (x_1, x_2) in comparison to the content in the absence of media; when $\overline{R} > 0$ (in pink), the media has increased the prevalence of content with the (ideology, quality) pair (x_1, x_2).

ideological position of 0.613636 and a quality of 0.3125 in the Media Bias Chart and is followed by about 28% of the nonmedia accounts) is very popular and near the final location of the conservative community in (ideology, quality) space.

VIII. CONCLUSIONS AND DISCUSSION

We introduced a model for the spread of content in an online social network, with both media and nonmedia nodes, that accounts for the effects of both ideology and content quality on the spreading dynamics. We based the spreading dynamics of our model on a bounded-confidence mechanism, such that accounts spread content that is sufficiently close to their current ideological position. In the most sophisticated version of our model, the quality of content determines how close an account's ideological position should be to the content's ideological position for the account to share the content. Therefore, low-quality content is shared only when it supports an account's existing ideological biases. As far as we are aware, our model is the first mathematical model to explicitly incorporate the effects of media quality on spreading dynamics. This is a key novelty of our work.

We conducted simulations for our content-spreading model for media content with both one and two ideological dimensions. Using results from our simulations, we quantified the level of media "entrainment" (i.e., how much the media affects the ideological positions of nonmedia accounts) for a variety of network architectures. We examined how the amount of media entrainment of the nonmedia accounts in a network depends on the numbers of media accounts and on the number of followers per media account in that network. We found that media impact increases when one increases the number of nonmedia accounts, the expected mean number of nonmedia accounts that such an account follows, or the receptiveness of such an account. We also observed an interesting relationship between media entrainment and convergence time in our model, with higher levels of media entrainment correlating positively with longer convergence times. Finally, we simulated a version of our model that accounts for content quality in the spreading dynamics. Using a hand-curated set of media inputs from real media outlets and their numbers of followers on Twitter, we demonstrated that this version of our model produces polarization in both ideology and quality of content. Specifically, it yields a community of high-quality content in the center of the political spectrum and a conservative community of low-quality content. Our model provides a useful step towards increasing understanding of how media content quality affects the spread of online content.

Our model is a simplistic model for the spread of media content in online social networks. This is a complex system, and naturally it is not appropriate to view our model as a perfectly accurate mathematical description of such phenomena. Instead, our model provides a starting point for exploring the mechanisms that contribute to content spreading dynamics and echo chambers. There are many worthwhile ways to generalize our model. For example, our assumption that accounts choose whether to spread content based on a universal confidence parameter c is a naive simplification, as is the homogeneity of nonmedia accounts in general. Spreading behavior surely depends on individual characteristics, as has been explored in models that include zealots [54] and in other models of social dynamics [24,55,56]. Augmenting our model of content spreading by incorporating account heterogeneity is important future work. In such efforts, we expect that it will be insightful to explore the effects of structural homophily (for example, as was explored using the DeGroot model in [57]). Individuals in social networks are free to choose which accounts to follow (and which accounts to stop following or never follow), and such choices are sometimes driven by the desire to follow accounts that have similar ideologies [4,5]. This can increase structural homophily and exaggerate echo chambers in networks [22,58], because edges are more likely to arise (and persist) between nodes with similar ideologies.

There are also other interesting avenues for extending our work. For example, when modeling content with multiple ideological dimensions, one can give different weights to different dimensions. One can also incorporate spreading through multiple types of social media by generalizing our model to multilayer networks [59] or develop models of media influence that go beyond pairwise interactions by generalizing it to hypergraphs or simplicial complexes [60]. It will also be interesting to extend recent work on bounded-confidence opinion dynamics that coevolve with network structure [61] to include the effects of media. It is also important to explore the influence of time-dependence in network architecture [62] on the dynamics of our content updating rule, as accounts can follow new accounts and unfollow accounts. One way to explore such phenomena is to incorporate network rewiring, as in adaptive voter models [63].

A key strength of our model is its flexibility, as one can formulate generalizations of it (such as ones that we just discussed) in a straightforward way. Furthermore, one can already readily use our model to study many interesting phenomena. For example, although we only examined a few case studies of media probability distributions in the present paper, one can draw the number of media accounts, the number of followers per account, and the ideologies and qualities of media content from any probability distribution—either synthetic ones or ones that are inferred from empirical data—and it is important to explore these ideas in future work. It is particularly desirable for media nodes in a content-spreading model to produce content that follows some distribution of ideologies and qualities, rather than always using the same parameter values for such features. It will also be very interesting to include media nodes with ideologies x_M that are inferred from data, such as through sentiment analysis of news stories on a given topic (as described in Ref. [19] and references therein). Using procedures such as topic modeling and sentiment analysis can generate a probability distribution of media accounts in ideological space as an output, which can then be fed into our model. There has been some recent progress in this direction. For example, Ye and Skiena [64] measured media bias and quality from the web pages and tweets of media sources, and Albanese *et al.* [65] created a two-dimensional model based on semantic analysis to quantify media influence on voting. Another important issue is the examination of transient dynamics of content spreading, as it is necessary to go beyond our focus on the properties of ideological positions after long times or at convergence.

Developing mathematical approaches for analyzing and quantifying the dynamics of content spreading has ramifications for how to mitigate the spread of undesired content and promote the spread of desired content in social media. One potential impact is the development of control strategies and "fake-news filters" (e.g., by flagging content that is below some threshold value on the quality axis in a model like ours) that are reminiscent of spam filters. Our work provides a step toward the development of novel algorithms to encourage selective spreading of high-quality or desirable content. Another area for which further development of such models may also be fruitful is in bot detection, where most existing algorithms rely on network measures, followership data, activity rates, or linguistic cues [7] (all of which are straightforward to manipulate by malicious actors), rather than using spreading dynamics. Advances in these modeling efforts will also yield insights into the theory of online content spread and help bridge the gap between simple spreading models and realistic investigations of spreading on social media.

ACKNOWLEDGMENTS

MAP was supported by the National Science Foundation (Grant No. 1922952) through the Algorithms for Threat Detection (ATD) program. We thank Franca Hoffmann, Alex Pan, and two anonymous referees for their helpful comments.

[1] Pew Research Center, News use across social media platforms 2018 (2018), available at http://www.journalism.org/2018/09/10/news-use-across-social-media-platforms-2018/.

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

[3] R. M. Bond, C. J. Fariss, J. J. Jones, A. D. Kramer, C. Marlow, J. E. Settle, and J. H. Fowler, A 61-million-person experiment in social influence and political mobilization, Nature (London) **489**, 295 (2012).

[4] E. Bakshy, S. Messing, and L. A. Adamic, Exposure to ideologically diverse news and opinion on Facebook, Science **348**, 1130 (2015).

[5] A. Bessi, G. Caldarelli, M. Del Vicario, A. Scala, and W. Quattrociocchi, Social determinants of content selection in the age of (mis)information, in *International Conference on Social Informatics*, SocInfo 2014 (Springer International Publishing, Cham, Switzerland, 2014), pp. 259–268.

[6] M. Cinelli, E. Brugnoli, A. L. Schmidt, F. Zollo, W. Quattrociocchi, and A. Scala, Selective exposure shapes the Facebook news diet, arXiv:1903.00699 (2019).

[7] E. Ferrara, O. Varol, C. Davis, F. Menczer, and A. Flammini, The rise of social bots, Commun. ACM **59**, 96 (2016).

[8] K. T. Poole, *Spatial Models of Parliamentary Voting* (Cambridge University Press, Cambridge, UK, 2005).

[9] K. T. Poole and H. L. Rosenthal, *Ideology and Congress* (Transaction Publishers, Piscataway, NJ, 2007).

[10] J. Lewis, K. Poole, H. Rosenthal, A. Boche, A. Rudkin, and L. Sonnet, Voteview: Congressional roll-call votes database (2019), available at http://voteview.com.

[11] L. Sirovich, A pattern analysis of the second Rehnquist U.S. Supreme Court, Proc. Natl. Acad. Sci. USA **100**, 7432 (2003).

[12] M. A. Porter, P. J. Mucha, M. E. J. Newman, and C. M. Warmbrand, A network analysis of committees in the U.S. House of Representatives, Proc. Natl. Acad. Sci. USA **102**, 7057 (2005).

[13] S. Vosoughi, D. Roy, and S. Aral, The spread of true and false news online, Science **359**, 1146 (2018).

[14] A. Guille, H. Hacid, C. Favre, and D. A. Zighed, Information diffusion in online social networks: A survey, ACM Sigmod Record **42**, 17 (2013).

[15] F. Jin, E. Dougherty, P. Saraf, Y. Cao, and N. Ramakrishnan, Epidemiological modeling of news and rumors on Twitter, in *Proceedings of the 7th Workshop on Social Network Mining and Analysis*, SNAKDD '13 (Association for Computing Machinery, New York, NY, USA, 2013), Article No. 8.

[16] Q. Liu, T. Li, and M. Sun, The analysis of an SEIR rumor propagation model on heterogeneous network, Physica A **469**, 372 (2017).

[17] J. Wang, L. Zhao, and R. Huang, SIRaRu rumor spreading model in complex networks, Physica A **398**, 43 (2014).

[18] L. Zhao, H. Cui, X. Qiu, X. Wang, and J. Wang, SIR rumor spreading model in the new media age, Physica A **392**, 995 (2013).

[19] B. Batrinca and P. C. Treleaven, Social media analytics: A survey of techniques, tools and platforms, AI & Society **30**, 89 (2015).

[20] C. R. Sunstein, *Republic.com 2.0* (Princeton University Press, Princeton, NJ, 2009).

[21] S. Flaxman, S. Goel, and J. M. Rao, Filter bubbles, echo chambers, and online news consumption, Public Opin. Q. **80**, 298 (2016).

[22] M. Del Vicario, G. Vivaldo, A. Bessi, F. Zollo, A. Scala, G. Caldarelli, and W. Quattrociocchi, Echo chambers: Emotional contagion and group polarization on Facebook, Sci. Rep. **6**, 37825 (2016).

[23] W. Cota, S. C. Ferreira, R. Pastor-Satorras, and M. Starnini, Quantifying echo chamber effects in information spreading over political communication networks, EPJ Data Sci. **8**, 35 (2019).

[24] *Complex Spreading Phenomena in Social Systems: Influence and Contagion in Real-World Social Networks*, edited by S. Lehmann and Y.-Y. Ahn (Springer International Publishing, Cham, Switzerland, 2018).

[25] J. D. O'Brien, I. K. Dassios, and J. P. Gleeson, Spreading of memes on multiplex networks, New J. Phys. **21**, 025001 (2019).

[26] F. Baumann, P. Lorenz-Spreen, I. M. Sokolov, and M. Starnini, Modeling Echo Chambers and Polarization Dynamics in Social Networks, Phys. Rev. Lett. **124**, 048301 (2020).

[27] U. Chitra and C. Musco, *Understanding filter bubbles and polarization in social networks*, WISDOM '19 (KDD 2019, August 5th, Anchorage, AK, USA), arXiv:1906.08772.

[28] H. A. Prasetya and T. Murata, Modeling the co-evolving polarization of opinion and news propagation structure in social media, in *International Workshop on Complex Networks and their Applications* (Springer International Publishing, Cham, Switzerland, 2018), pp. 314–326.

[29] A. C. R. Martins, Mobility and social network effects on extremist opinions, Phys. Rev. E **78**, 036104 (2008).

[30] Y.-l. Chuang and M. R. D'Orsogna, Mathematical models of radicalization and terrorism, arXiv:1903.08485.

[31] H. Noorazar, K. R. Vixie, A. Talebanpour, and Y. Hu, From classical to modern opinion dynamics, arXiv:1909.12089.

[32] G. Deffuant, D. Neau, F. Amblard, and G. Weisbuch, Mixing beliefs among interacting agents, Adv. Complex Syst. **3**, 87 (2000).

[33] R. Hegselmann and U. Krause, Opinion dynamics and bounded confidence: Models, analysis, and simulation, J. Artif. Soc. Soc. Simul. **5**, 3 (2002).

[34] D. Bhat and S. Redner, Opinion formation under antagonistic influences, Phys. Rev. E **100**, 050301 (2019).

[35] D. Bhat and S. Redner, Polarization and consensus by opposing external sources, J. Stat. Mech. (2020) 013402.

[36] M. Mobilia, A. Petersen, and S. Redner, On the role of zealotry in the voter model, J. Stat. Mech. (2007) P08029.

[37] D. Acemoğlu, G. Como, F. Fagnani, and A. Ozdaglar, Opinion fluctuations and disagreement in social networks, Math. Oper. Rese. **38**, 1 (2013).

[38] Y. Yi and S. Patterson, Disagreement and polarization in two-party social networks, arXiv:1911.11338.

[39] R. Hegselmann and U. Krause, Opinion dynamics under the influence of radical groups, charismatic leaders, and other constant signals: A simple unifying model, Netw. Heterog. Media, **10**, 477 (2015).

[40] J. Lorenz, Fostering consensus in multidimensional continuous opinion dynamics under bounded confidence, in *Managing Complexity: Insights, Concepts, Applications*, edited by D. Helbing, Understanding Complex Systems (Springer-Verlag, Heidelberg, Germany, 2008), pp. 321–334.

[41] J. Li and R. Xiao, Agent-based modelling approach for multidimensional opinion polarization in collective behaviour, J. Artif. Soc. Soc. Simul. **20**, 2 (2017).

[42] S. T. Stamoulas and M. Rathinam, Convergence, stability, and robustness of multidimensional opinion dynamics in continuous time, SIAM J. Control Optim. **56**, 1938 (2018).

[43] F. Bullo, *Lectures on Network Systems*, 1.3 ed. (Kindle Direct Publishing, 2019) with contributions by Jorge Cortés, Florian Dörfler, and Sonia Martínez, http://motion.me.ucsb.edu/book-lns/.

[44] P. Sobkowicz, Extremism without extremists: Deffuant model with emotions, Front. Phys. **3**, 17 (2015).

[45] 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).

[46] L. Adamic, T. Lento, E. Adar, and P. Ng, The evolution of memes on Facebook (2014), available at https://www.facebook.com/notes/facebook-data-science/the-evolution-of-memes-on-facebook/10151988334203859.

[47] M. H. DeGroot, Reaching a consensus, J. Am. Stat. Assoc. **69**, 118 (1974).

[48] W. Mei, F. Bullo, G. Chen, and F. Dörfler, Occam's razor in opinion dynamics: The weighted-median influence process, arXiv:1909.06474.

[49] P. Jia, A. MirTabatabaei, N. E. Friedkin, and F. Bullo, Opinion dynamics and the evolution of social power in influence networks, SIAM Rev. **57**, 367 (2015).

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

[51] S. Lehmann and P. Sapieżyński, You're here because of a robot (2013), available at https://sunelehmann.com/2013/12/04/youre-here-because-of-a-robot/.

[52] D. J. Watts and S. H. Strogatz, Collective dynamics of 'small-world' networks, Nature (London) **393**, 440 (1998).

[53] V. Otero, The Media Bias Chart project (2019), available at http://www.adfontesmedia.com/intro-to-the-media-bias-chart/.

[54] A. Waagen, G. Verma, K. Chan, A. Swami, and R. D'Souza, Effect of zealotry in high-dimensional opinion dynamics models, Phys. Rev. E **91**, 022811 (2015).

[55] M. A. Porter and J. P. Gleeson, *Dynamical Systems on Networks: A Tutorial*, Frontiers in Applied Dynamical Systems: Reviews and Tutorials Vol. 4 (Springer International Publishing, Cham, Switzerland, 2016).

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

[57] P. Dandekar, A. Goel, and D. T. Lee, Biased assimilation, homophily, and the dynamics of polarization, Proc. Natl. Acad. Sci. USA **110**, 5791 (2013).

[58] A. Bessi, F. Petroni, M. Del Vicario, F. Zollo, A. Anagnostopoulos, A. Scala, G. Caldarelli, and W. Quattrociocchi, Homophily and polarization in the age of misinformation, Eur. Phys. J.: Spec. Top. **225**, 2047 (2016).

[59] M. Kivelä, A. Arenas, M. Barthelemy, J. P. Gleeson, Y. Moreno, and M. A. Porter, Multilayer networks, J. Complex Networks **2**, 203 (2014).

[60] R. Lambiotte, M. Rosvall, and I. Scholtes, From networks to optimal higher-order models of complex systems, Nat. Phys. **15**, 313 (2019).

[61] M. Brede, How does active participation effect consensus: Adaptive network model of opinion dynamics and influence maximizing rewiring, Complexity **2019**, 1486909 (2019).

[62] P. Holme and J. Saramäki, Temporal networks, Phys. Rep. **519**, 97 (2012).

[63] R. Durrett, J. P. Gleeson, A. L. Lloyd, P. J. Mucha, F. Shi, D. Sivakoff, J. E. Socolar, and C. Varghese, Graph fission in an evolving voter model, Proc. Natl. Acad. Sci. USA **109**, 3682 (2012).

[64] J. Ye and S. Skiena, MediaRank: Computational ranking of online news sources, in *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, KDD 2019 (Association for Computing Machinery, New York, NY, USA, 2019), pp. 2469–2477.

[65] F. Albanese, C. J. Tessone, V. Semeshenko, and P. Balenzuela, A data-driven model for mass media influence in electoral context, arXiv:1909.10554.