Information Sciences 593 (2022) 49–63

Contents lists available at ScienceDirect

Information Sciences

journal homepage: www.elsevier.com/locate/ins

Biased opinion dynamics: when the devil is in the details ☆

Aris Anagnostopoulos [a], Luca Becchetti [a], Emilio Cruciani [b,*], Francesco Pasquale [c], Sara Rizzo [d]

[a] Sapienza University of Rome, Rome, Italy
[b] Paris-Lodron University of Salzburg, Salzburg, Austria
[c] Tor Vergata University of Rome, Rome, Italy
[d] Gran Sasso Science Institute, L'Aquila, Italy

ARTICLE INFO

Article history:
Received 2 August 2021
Received in revised form 10 December 2021
Accepted 30 January 2022
Available online 2 February 2022

Keywords:
Opinion dynamics
Majority dynamics
Voter model
Social networks
Consensus
Markov chains

ABSTRACT

We study opinion dynamics in multi-agent networks when a bias toward one of two possible opinions exists, for example reflecting a status quo versus a superior alternative. Our aim is to investigate the combined effect of bias, network structure, and opinion dynamics on the convergence of the system of agents as a whole. Models of such evolving processes can easily become analytically intractable. In this paper, we consider a simple yet mathematically rich setting, in which all agents initially share an initial opinion representing the status quo. The system evolves in steps. In each step, one agent selected uniformly at random follows an underlying update rule to revise its opinion on the basis of those held by its neighbors, but with a probabilistic bias towards the superior alternative. We analyze convergence of the resulting process under well-known update rules. The framework we propose is simple and modular, but at the same time complex enough to highlight a nonobvious interplay between topology and underlying update rule.

1. Introduction

Opinion formation in social groups has been the focus of extensive research. Even though many models considered in the literature confer the same *intrinsic* value to all opinions [13], one might expect a group to quickly reach consensus on a clearly superior alternative, if present. Yet, phenomena such as *groupthink* may delay or even prevent such an outcome, a phenomenon that has been noted in the business world.[1,2] This phenomenon has even been considered as a factor behind France's defeat by Germany in the second world war [1].

In this perspective, we investigate models of opinion formation in which a bias towards one of two possible opinions exists, for instance, reflecting intrinsic superiority of one alternative over the other.[3] In particular, our aim is to investigate in a simple and mathematically tractable way the effect of the bias and the network structure on the time it takes for the system of agents as a whole to converge to the superior opinion. In the remainder, we use labels 0 and 1 for the two opinions and we assume that 1 is the *dominant* opinion, that is, the one towards which the agents have a bias. We investigate this question in a mathematically tractable setting, informally described as follows.

An extended abstract appeared in the Proceedings of IJCAI 2020 [2].
* Corresponding author at: Jakob-Haringer-Straße 2, 5020 Salzburg, Austria.
E-mail address: emilio.cruciani@sbg.ac.at (E. Cruciani).
[1] https://smallbusiness.chron.com/can-groupthink-affect-organization-26044.html
[2] https://www.linkedin.com/pulse/groupthink-curse-innovation-business-transformation-fisher-chejoski
[3] Characterizing the notion of "superiority" is typically context-dependent and may be far from obvious.

https://doi.org/10.1016/j.ins.2022.01.072
0020-0255/© 2022 The Author(s). Published by Elsevier Inc.

Assume some underlying opinion dynamics $\mathcal{D}$. Starting from an initial state in which all agents share opinion 0, the system evolves in discrete time steps. In each time step, one agent is selected uniformly at random. With some probability α, the agent adopts 1, and with probability $1-\alpha$, the agent follows $\mathcal{D}$ to revise its opinion on the basis of those held by its neighbors in an underlying network.

Although the general model we consider is simple and, under mild conditions on $\mathcal{D}$, the family of processes it describes always admits global adoption of opinion 1 as the only absorbing state, convergence to this absorbing state exhibits a rich variety of behaviors, which depends on the interplay between the network structure and the underlying opinion dynamics in nonobvious ways. The relatively simple, yet general, model we consider allows analytical investigation of the following question:

> How does a particular combination of network structure and opinion dynamics affect convergence to global adoption of the dominant opinion? In particular, how conducive is a particular combination to rapid adoption?.

Even though results implying a negative effect of the network density on the convergence time have been proposed in the past, albeit for quite different models (e.g., [34]), our findings suggest that there might be more to the issue. In particular, the interplay between opinion dynamics and underlying network structure seems more complex than anticipated, with the former playing a key role in amplifying network effects.

In this respect, it is important to note that a number of continuous-time models of (discrete) opinion formation processes have been proposed in the past (see Section 2). Having their roots in the study of compartmental models of spreading dynamics [38], these models can often be richer than the ones we consider in this paper. In particular, they afford empirical investigation of a number of interesting phenomena, typically in regimes in which network size tends to become very large and correlations can to some degree be neglected. Despite being very useful, these models tend to be extremely complex. The simplifying assumptions introduced to make them tractable, often prevent a fine-grained, mathematically rigorous investigation of the relationships between opinion formation model, underlying network structure, and bias. This problem is also acknowledged by recent work in this area (Section I, [39]): "we still lack a general theoretical framework systematically linking models of information spreading, network structure, and algorithmic bias." In contrast, our goal here is to trade some expressivity of the model for the ability to rigorously characterize it, while hopefully retaining enough complexity to capture basic, yet fundamental phenomena at work.

Organization of the paper and main results. We discuss work that is more or less closely related to the topic of this study in Section 2, and we present and formalize the general model that we consider in Section 3.

In Section 4, we show that the expected time for consensus on the dominant opinion grows exponentially with the minimum degree under the *majority* update rule, in which agents update their opinion to the majority opinion in their neighborhoods [29]. Using asymptotic notation and denoting the number of agents in the network by n, we obtain that the convergence time is super-polynomial in expectation whenever the minimum degree is $\omega(\log n)$. One might wonder whether the converse occurs, namely, whether a logarithmic maximum degree admits an (expected) polynomial convergence to the absorbing state. Even though we prove that this is indeed the case for specific topologies such as cycles or restricted graph families, this does not seem to hold in general (see discussion in Section 7).

The results for the majority rule discussed in the previous paragraph are at odds with those we obtain in Section 5 for the *voter model*, in which agents copy the opinion of a randomly selected neighbor [31]. In this case, convergence to the absorbing state occurs within $\mathcal{O}\left(\frac{1}{\alpha} n \log n\right)$ time steps *with high probability*, regardless of the underlying network structure.

Empirical analysis presented in Section 6 provides results that are consistent with our theoretical findings. In particular, the behavior predicted for the voter model closely matches empirical results. For majority, empirical analysis complements our theoretical findings, highlighting a sharp phase transition in absorption times, with a threshold that depends on the bias and the graph's degree distribution in a nonobvious fashion.

Finally, in Section 7, we discuss generalizations of our framework and possible directions for further investigation.

To summarize, this study provides a simple mathematical framework to investigate the interplay between opinion dynamics and underlying network structure in a unified setting, allowing comparison of different update rules with respect to a common framework. In this respect, we hope that our work moves in the direction of a shared framework to investigate opinion dynamics, as advocated in [13], complementing the aforementioned more complex, but harder to analyze, models. The rigorous analysis of the simpler models of this paper, helps us better understand, at least qualitatively, important driving forces behind the dynamics of agent interactions in the presence of bias.

2. Related Work

The problem we consider touches a number of areas where similar settings have been considered, with various motivations. The corresponding literature is vast and providing an exhaustive review is unfeasible here. In the paragraphs that follow, we discuss contributions that most closely relate to the topic of this paper.

Opinion diffusion and consensus. Opinion dynamics are widely used to investigate how groups of agents modify their beliefs under the influence of other agents and possibly exogenous factors. A number of models have been proposed in the more or less recent past, mostly motivated by phenomena that arise in several areas, ranging from social sciences to physics and biology. Refer to [13] and references therein for a recent, general overview of opinion dynamics in multi-agent systems. A first distinction is between settings in which the set of possible beliefs is continuous, for instance, the interval $[0, 1]$. This setting has been the focus of extensive research in social sciences and economics [3,17,21,22].

In this paper, we consider the case in which opinions are drawn from a discrete set, a setting that also received significant attention in the recent past. In particular, we focus on the *majority rule* and the *voter model*. Investigation of the majority update rule originates from the study of agreement phenomena in spin systems [29], whereas the voter model was motivated by the study of spatial conflict between species in biology and interacting stochastic processes and particle systems in probability theory and statistics [12,26,31]. These two models have received renewed attention in the recent past, the focus mostly being on the time required to achieve consensus or conditions under which they achieve consensus on one of the initial opinions with a minimum degree of confidence. The voter model is by now well understood. In particular, increasingly tight bounds on convergence time for general and specific topologies have been proposed over the recent past [15,25], and it is known that the probability of one particular opinion to prevail is proportional to the sum of the degrees of nodes holding that opinion at the onset of the process [18]. In all these works there is not a bias toward one of the alternatives, which is the topic of this paper.

Consensus and network structure. Network structure has been known to play an important role in opinion diffusion and influence spreading for quite some time [35], under a variety of models. For example, multiple works have studied the consensus under the voter model and the dependence of its convergence on the underlying network topology [15,18,25]. For the majority dynamics, Auletta et al. [5] characterized the topologies for which an initial majority can be subverted, showing that this is possible for all but a handful of topologies, including cliques and quasi-cliques. On the other hand, regardless of the network, there is always an initial opinion distribution such that the final majority will reflect the initial one; however, computing an initial opinion configuration that will subvert an initial majority is topology-dependent and NP-hard in general [8].

A number of recent contributions have investigated (among other aspects) the relationship between network structure and consensus in opinion formation games [19,20], whereas Auletta et al. [7] studied extensions of the Friedkin-Johnsen model to evolving networks.

Even though a high expansion of the underlying graph typically implies fast convergence [14,27] in many opinion dynamics, some recent work explicitly points to potentially adverse effects of network structure on the spread of innovation, at least in scenarios where the opinion update occurs on the basis of private utilities that reflect both the degree of local consensus and the intrinsic value of the competing opinions [34,41].

Some of our findings are qualitatively consistent with previous work albeit under completely different models (in particular, [34]). However, our overall approach is very different, because it completely decouples the mechanism of opinion formation from modelling of the bias, allowing for a clear-cut mathematical characterization of the interplay between bias, underlying opinion dynamics, and network structure.

Follow-up work to a preliminary version of this paper [2] investigated the role of small sets of high-degree nodes (or elites) to influence the outcome of the majority dynamics. In particular, Out and Zehmakan [37] identified the addition of random connections and/or some of degree bias towards the currently held opinion as effective strategies to mitigate the ability of elites to drive consensus towards a predefined opinion.

Forms of bias different from ours. Bias in opinion dynamics has been considered previously in the literature. We briefly review contributions that are at least loosely related to our framework. For both voter and majority update rules, Mukhopadhyay et al. [36] introduced bias in the form of different, opinion-dependent firing rate frequencies of the Poisson clocks that trigger agents' opinion updates, implicitly enforcing a bias toward the opinion with lower associated rate. Despite being different, their model is similar to ours in spirit and some of their results for the voter model are consistent with ours. Yet, these results only apply in expectation and to very dense networks with degree $\Omega(n)$, whereas our results for the voter model hold for every undirected graph. More recently, Bahrani et al. [9] investigated the behavior of majority dynamics when agents' opinions are initialized to 0 (incorrect) or 1 (correct) with probabilities $1/2 - \delta$ and $1/2 + \delta$ respectively. The authors prove high probability of convergence to a majority of correct opinions within $\mathcal{O}(n \log n / \log \log n)$ steps.

A form of bias in the communication channels between the agents is considered by Cruciani et al. [16], where one of the two opinions is always transmitted correctly to the neighbors, whereas there is a fixed probability of incorrectly transmitting the other one. The results for both voter and majority models are coherent with the ones presented in this paper, even if in the work of Cruciani et al. they hold for a synchronous setting, where in each round all nodes update simultaneously.

A somewhat related line of research addresses the presence of stubborn agents and of zealots. Loosely speaking, stubborn agents have a bias toward some (initially or currently) held opinion. Zealots are agents who never deflect from some initial opinion. Restricting to the discrete-opinion setting, which is the focus of this paper,[4] the role of zealots and their ability to subvert an initial majority have been investigated for the voter model (see [33] and follow-up work), and Auletta et al. [6] investigated majority dynamics in the presence of stubborn agents that are biased toward the currently held opinion, providing a full characterization of conditions under which an initial majority can be subverted.

Continuous-time dynamics. Continuous-time, discrete-opinion models have been proposed in the past and form an active research area, following a long-standing tradition that is rooted in the analysis of epidemic processes [38]. We only mention a few contributions that are representative of this line of research. These approaches generally rely on mean-field or finer-grained, continuous approximations [23,24] in which, roughly speaking, the probabilities that characterize discrete models

[4] For the continuous case, there is a vast literature; see the seminal paper [22] and follow-up work.

correspond to transition rates between different opinions across the underlying population of agents. For example, [23,24] consider a refined mean-field analysis of a number of spreading/opinion diffusion dynamics, including a noisy variant of majority, whereas a continuous, noisy voter model has been considered recently [11], under the assumption of lack of correlation among network edges. More recent work addresses the issue of bias in the opinion formation process. For example, Peralta et al. [39] model the presence of (algorithmic) bias by a single parameter, a so-called "bias intensity," which modifies the transition rates defined in [23,24] in a simple way.

Continuous-time models can be rich and afford empirical investigation of a number of interesting phenomena. The resulting systems of differential equations can be accurate approximations of some original (continuous or discrete) model. They often allow to simulate the evolution of models that are generally more complex than those that are analytically tractable. Sometimes, these models also allow to derive quantitative predictions on parameters of interest, which empirically prove accurate under suitable conditions, for example, when the correlations that can be present in the system within sub-populations of the agents can be considered negligible. On the negative side, attempts at a theoretical analysis often require strong simplifying assumptions to ensure tractability, which in turn can make it harder (when not impossible) to derive rigorous, analytical results. Also, these models can prove inaccurate with respect to the original ones in some cases. For example, mean-field models can behave poorly on sparse networks, or on networks that are close to exhibiting properties that induce a transition in the behavior of the underlying dinamics [23]. As a result, deriving quantitative, mathematically rigorous characterizations of the complex relationships involving opinion dynamics, bias, and network structure, can become unfeasible in such settings.

3. Notation and Preliminaries

Let $G = (V, E)$ be an undirected graph with $|V| = n$ nodes, each representing an agent. Without loss of generality, we assume that $V = [n] := \{1, \dots, n\}$. The system evolves in discrete time steps[5] and, at every time step $t \in \mathbb{N}$, each node $v \in V$ holds an *opinion* $x_v^{(t)} \in \{0, 1\}$. We use the term opinion liberally here, in the sense that 0 and 1 in general represent competing alternatives, whose meaning is context-dependent. We denote by $\mathbf{x}^{(t)} = \left(x_1^{(t)}, \dots, x_n^{(t)}\right)^\intercal$ the corresponding *state* of the system at time t. We assume that the initial state of the system is $\mathbf{x}^{(0)} = \mathbf{0} = (0, \dots, 0)^\intercal$; this assumption is discussed in Section 7. For each $v \in V$, we denote the neighborhood of v by $N_v := \{u \in V : \{u, v\} \in E\}$ and the degree of v by $d_v := |N_v|$. Finally, $\Delta := \min_{v \in V} d_v$ is the minimum degree among the nodes in G.

Our framework assumes that the agents exhibit a bias toward one of the opinions (e.g., reflecting intrinsic superiority of a technological innovation over the status quo), without loss of generality 1, which we henceforth call the *dominant opinion*. We model bias as a probability, with a parameter $\alpha \in (0, 1]$. All dynamics we consider are *Markovian*, that is, given the underlying graph G, the distribution of the state $\mathbf{x}^{(t)}$ at time t only depends on the state $\mathbf{x}^{(t-1)}$ at the end of the previous time step. Moreover, they have $\mathbf{x} = \mathbf{1} = (1, \dots, 1)^\intercal$ as the only absorbing state. We use τ to denote the *absorption time*, which is the number of time steps for the process to reach the absorbing state $\mathbf{1}$. Finally, for a family of events $\{\mathcal{E}_n\}_{n \in \mathbb{N}}$, we say that $\mathcal{E}_n$ occurs *with high probability* (*w.h.p.*, in short), if a constant $\gamma > 0$ exists such that $\mathbf{P}(\mathcal{E}_n) = 1 - \mathcal{O}(n^{-\gamma})$, for every sufficiently large n.

4. Absorption Time for Majority Dynamics

In this section, we investigate the time to reach consensus on the dominant opinion under the majority update rule. More formally, we study the following random process: Starting from the initial state $\mathbf{x}^{(0)} = (0, \dots, 0)^\intercal$, in each time t a node $u \in [n]$ is chosen uniformly at random and u updates its value according to the rule

$$x_u^{(t)} = \begin{cases} 1 & \text{with probability } \alpha, \\ M_G(u, \mathbf{x}) & \text{with probability } 1 - \alpha, \end{cases}$$

where $\alpha \in (0, 1]$ is the bias toward the dominant opinion 1 and $M_G(u, \mathbf{x})$ is the value held in configuration $\mathbf{x}^{(t-1)} = \mathbf{x}$ by the majority of the neighbors of node u in graph G:

$$M_G(u, \mathbf{x}) = \begin{cases} 0 & \text{if } \sum_{v \in N_u} x_v < |N_u|/2, \\ 1 & \text{if } \sum_{v \in N_u} x_v > |N_u|/2, \end{cases}$$

and ties are broken uniformly at random, that is, if $\sum_{v \in N_u} x_v = |N_u|/2$ then $M_G(u, \mathbf{x}) = 0$ or 1 with probability 1/2, independently of all the other random choices of this process.

It is straightforward to see that for every positive α, the above Markov chain has $\mathbf{1}$ as the only absorbing state. However, the rate of convergence is strongly influenced by the underlying graph G. In Section 4.1 we prove a lower bound on the

[5] This is equivalent to the asynchronous model in which a node revises its opinion at the arrival of an independent Poisson clock with rate 1 [10].

expected absorption time, which depends exponentially on the minimum degree. This result implies super-polynomial expected absorption times for graphs whose minimum degree is $\omega(\log n)$. On the other hand, in Section 4.2 we prove that the absorption time is $\mathcal{O}(n\log n)$ on cycle graphs, and in Section 4.3 we briefly discuss about further graph families with sub-logarithmic maximum degree and polynomial (expected) absorption time.

4.1. Convergence on High-Density Graphs

In this section we prove a general lower bound on the expected absorption time, which only depends on the minimum degree Δ. To this purpose, we use the following standard lemma on birth-and-death chains[6] (see, for example, Section 17.3 [30] for a proof).

Lemma 4.1. *Let $\{X_t\}_t$ be a birth-and-death chain with state space $\{0,1,\dots,n\}$, such that for every $1 \leqslant k \leqslant n-1$*

$$\mathbf{P}(X_{t+1}=k+1\,|\,X_t=k)=p,$$
$$\mathbf{P}(X_{t+1}=k-1\,|\,X_t=k)=q,$$
$$\mathbf{P}(X_{t+1}=k\,|\,X_t=k)=r,$$

with $p+q+r=1$. For every $i\in\{0,1,\dots,n\}$ let τ_i be the first time the chain hits state i, that is, $\tau_i=\min\{t\,|\,X_t=i\}$. If $0<p<1/2$, the probability that starting from state k the chain hits state n before state 0 is

$$\mathbf{P}_k(\tau_n<\tau_0)=\frac{(q/p)^k-1}{(q/p)^n-1}\leqslant\left(\frac{p}{q}\right)^{n-k}.$$

First we show that, for $\alpha \geqslant 1/2+\varepsilon$, every graph with minimum degree $\Delta=\Omega\left(\frac{\log n}{\varepsilon^2}\right)$ has $\mathcal{O}(n\log n)$ absorption time, w.h.p.

Lemma 4.2 (*Fast convergence for $\alpha>1/2$*). *Assume $\varepsilon>0$ and $\gamma>0$ are arbitrarily small constants. Then, if $\alpha>1/2+\varepsilon$, every graph with minimum degree at least $(1+\gamma)\frac{1+2\varepsilon}{\varepsilon^2}\ln n$ has absorption time at most $2(1+\gamma)n\ln n$, with probability at least $1-\frac{3}{n^\gamma}$.*

Proof. The proof proceed in two steps. We first prove that, w.h.p., within $\mathcal{O}(n\log n)$ time steps, every node has revised its opinion at least once. Let T denote the first time at which every node has revised its opinion at least once. Because every node is chosen independently with probability $1/n$ in every time step, this is a case of the coupon collector problem. Indeed, considering any particular node u, the probability that u is never sampled in any of the first $(1+\gamma)n\ln n$ time steps is $\left(1-\frac{1}{n}\right)^{(1+\gamma)n\ln n}$. Hence, using a union bound:

$$\mathbf{P}(T>(1+\gamma)n\ln n)\leqslant n\left(1-\frac{1}{n}\right)^{(1+\gamma)n\ln n}\leqslant ne^{-(1+\gamma)\ln n}=\frac{1}{n^\gamma}.$$

For the second step, consider again the first time T in which every node has revised its opinion at least once. Moreover, for every node u, denote by $T(u)$ the last time step prior to T in which u revised its opinion. Next, for every node u we define a binary variable X_u, such that $X_u=1$ if u adopts 1 at time $T(u)$, and $X_u=0$ otherwise. We have $\mathbf{P}(X_u=1)\geqslant\alpha>\frac{1}{2}+\varepsilon$, so that, for every node u, the expected number of u's neighbors that have adopted opinion 1 at time T is $\mathbf{E}\left[\sum_{v\in N_u}X_v\right]>\left(\frac{1}{2}+\varepsilon\right)|N_u|$. Moreover, although the X_u's are in general dependent on each other, for every u we have $\mathbf{P}(X_u=1)>\frac{1}{2}+\varepsilon$, regardless of other nodes. This means that every X_u stochastically dominates an independent Bernoulli variable Z_u with parameter $\frac{1}{2}+\varepsilon$. Moreover, if $Z=\sum_{v\in N_u}Z_v$, we have $\mathbf{E}[Z]=\left(\frac{1}{2}+\varepsilon\right)|N_u|$. We therefore have, for every u:

$$\begin{aligned}\mathbf{P}\left(\sum_{v\in N_u}X_v\leqslant\tfrac{1}{2}|N_u|\right)&\leqslant\mathbf{P}\left(Z\leqslant\tfrac{1}{2}|N_u|\right)\\&=\mathbf{P}\left(Z\leqslant\left(1-\tfrac{2\varepsilon}{1+2\varepsilon}\right)\mathbf{E}[Z]\right)\\&\overset{(a)}{\leqslant}e^{-\frac{2\varepsilon^2}{(1+2\varepsilon)^2}\mathbf{E}[Z]}\\&=e^{-\frac{\varepsilon^2}{1+2\varepsilon}|N_u|}\\&\overset{(b)}{\leqslant}e^{-\frac{\varepsilon^2}{1+2\varepsilon}\cdot(1+\gamma)\frac{1+2\varepsilon}{\varepsilon^2}\ln n}\\&=e^{-(1+\gamma)\ln n}\\&=\frac{1}{n^{1+\gamma}},\end{aligned}$$

[6] Birth-and-death chains are Markov processes for which, from state k, a transition can only go to either state $k+1$ or state $k-1$.

where (a) follows from a straightforward application of Chernoff's bound to the sum of the Z_u's, and (b) holds because $\min_u |N_u| \geqslant (1+\gamma)\frac{1+2\varepsilon}{\varepsilon^2}\ln n$ by the lemma hypothesis. Finally, taking a union bound over all nodes in V we obtain:

$$\mathbf{P}\left(\exists u \in V : \sum_{v \in N_u} X_v \leqslant \frac{1}{2}|N_u|\right) \leqslant \frac{1}{n^{\gamma}}.$$

The arguments above allow us to conclude that w.h.p., in the very first time T in which each node has revised its opinion at least once, namely at time $T \leqslant (1+\gamma)n\ln n$, w.h.p., every node is surrounded by a majority of neighbors that adopted opinion 1. Using again a coupon collector argument, we conclude that, with probability at least $1 - 1/n^{\gamma}$, every node will revise its opinion at least one more time within another $(1+\gamma)n\ln n$ time steps, this time adopting opinion 1 with probability 1. Overall, consensus to opinion 1 is achieved with probability at least $1 - 3/n^{\gamma}$ within $2(1+\gamma)n\ln n$ time steps.

In the next theorem we prove that, as soon as α is smaller than 1/2, the absorption time instead becomes exponential in the minimum degree.

Theorem 4.3 (*Slow convergence for $\alpha < 1/2$*). *Let $G = (V, E)$ be an undirected graph with minimum degree Δ. Assume $\alpha \leqslant \frac{(1-\varepsilon)}{2}$, for an arbitrary constant $0 < \varepsilon < 1$. The expected absorption time for the biased opinion dynamics under the majority update rule is*

$$\mathbf{E}[\tau] \geqslant \frac{e^{\frac{\varepsilon^2}{6}\Delta}}{6n}.$$

Proof. Let $S^{(t)}$ be the set of nodes with value 1 at time t. For each node $u \in V$, let $n_u^{(t)}$ be the fraction of its neighbors with value 1 at time t:

$$n_u^{(t)} = \frac{|N_u \cap S^{(t)}|}{|N_u|}.$$

Finally, let $\bar{\tau}$ be the first time step in which $n_u^{(t)} \geqslant 1/2$ for at least one node $v \in V$, namely,

$$\bar{\tau} = \min\{t \in \mathbb{N} : n_u^{(t)} \geqslant 1/2, \text{ for some } u \in [n]\}.$$

Note that for each time $t \leqslant \bar{\tau}$ all nodes have a majority of neighbors sharing opinion 0, thus the selected agent at time t updates its state to 1 with probability α and to 0 with probability $1-\alpha$. Moreover, clearly $\tau \geqslant \bar{\tau}$. We next prove that $\mathbf{E}[\bar{\tau}] \geqslant e^{\frac{\varepsilon^2}{6}\Delta}/(6n)$, which implies our thesis.

Observe that, for a node u with degree d_u that has k neighbors with value 1 at some time and for every $t \leqslant \bar{\tau}$, the probabilities $p_k(u)$ and $q_k(u)$ of increasing and decreasing, respectively, the number of its neighbors by 1 are

$$p_k(u) = \frac{d_u - k}{n}\alpha, \text{ and } q_k(u) = \frac{k}{n}(1-\alpha).$$

Hence, because $\alpha \leqslant (1-\varepsilon)/2$, for every $k \geqslant d_u/(2+\varepsilon)$ we have that

$$\frac{p_k(u)}{q_k(u)} = \frac{d_u - k}{k}\cdot\frac{\alpha}{1-\alpha} \leqslant (1+\varepsilon)\cdot\frac{1-\varepsilon}{1+\varepsilon} = 1-\varepsilon.$$

Note that

$$\frac{d_u}{2} - \frac{d_u}{2+\varepsilon} = d_u\frac{\varepsilon}{2(2+\varepsilon)} \geqslant \frac{\varepsilon}{6}d_u.$$

From Lemma 4.1 it thus follows that, for each node u, as soon as the number of its neighbors with value 1 enters in the range $(d_u/(2+\varepsilon), d_u/2)$, the probability that it will reach $d_u/2$ before going back to $d_u/(2+\varepsilon)$ is at most

$$(1-\varepsilon)^{\varepsilon d_u/6} \leqslant e^{-\varepsilon^2 d_u/6} \leqslant e^{-\frac{\varepsilon^2}{6}\Delta},$$

using $(1-x)^x \leqslant e^{-x^2}$ for $x \in [0,1]$. Hence, if we denote by Y_u the random variable indicating the number of trials before having at least half of the neighbors of u at 1, we have that for every $t \geqslant 0$

$$\mathbf{P}(Y_u \geqslant t) \geqslant \left(1 - e^{-\frac{\varepsilon^2}{6}\Delta}\right)^t \geqslant e^{-(3t/2)e^{-\frac{\varepsilon^2}{6}\Delta}},$$

where in the last inequality we used that $1 - x \geqslant e^{-3x/2}$ for every $x \in [0, \frac{1}{2})$. Thus,

$$\mathbf{P}(Y_u < t) \leqslant 1 - e^{-(3t/2)e^{-\frac{\varepsilon^2}{6}\Delta}} \leqslant \frac{3t}{2}e^{-\frac{\varepsilon^2}{6}\Delta},$$

using $1 - e^{-x} \leqslant x$ for every x. Finally, by using the union bound over all nodes, we have that

$$\mathbf{P}(\bar{\tau} < t) = \mathbf{P}(\exists u \in [n] : Y_u < t) \leqslant n \cdot \frac{3t}{2} e^{-\frac{\varepsilon^2}{6}\Delta}.$$

Thus, for $\bar{t} = e^{\frac{\varepsilon^2}{6}\Delta}/3n$ we have $\mathbf{P}(\bar{\tau} < \bar{t}) \leqslant 1/2$ and the theorem follows from Markov's inequality:

$$\mathbf{E}[\bar{\tau}] \geqslant \bar{t}\,\mathbf{P}(\bar{\tau} \geqslant \bar{t}) \geqslant \frac{\bar{t}}{2}.$$

4.2. Fast Convergence on the Cycle

In this section, we prove that the absorption time on an n-node cycle graph is $\mathcal{O}\left(\frac{1}{\alpha} n \log n\right)$, w.h.p. We make use of the following *structural* lemma.

Lemma 4.4 (*Structural property of cycles*). *Let C_n be the cycle of n nodes and let every node $v \in V$ have an associated state $x_v \in \{0, 1\}$. Let us call B_i and S_i the set of nodes in state i such that: every node $v \in B_i$ has both neighbors in the opposite state and every node $v \in S_i$ has one single neighbor in the opposite state. The following holds:*

$$|B_0| + \frac{|S_0|}{2} = |B_1| + \frac{|S_1|}{2}.$$

Proof. Given any possible binary coloring of C_n each node v belongs to one of the following categories:

- $v \in B_i$: node v is in state i and both its neighbors are in state $j \neq i$.
- $v \in R_i$: node v is in state i, its left neighbor is in state i, and its right neighbor is in state $j \neq i$.
- $v \in L_i$: node v is in state i, its right neighbor is in state i, and its left neighbor is in state $j \neq i$.
- $v \in Z_i$: node v is in state i and zero of its neighbors are in state $j \neq i$, that is, both are in state i.

We also call $S_i = R_i \cup L_i$. Fig. 1 illustrates the eight (counting symmetries) possible categories. Let us consider a clockwise walk through C_n that returns to its starting point. Keeping into account the categories of the nodes previously described it is possible to generate a graph H_C that describes all possible binary configurations of a C_n graph, for every $n \in \mathbb{N}$. We call H_C the *cycle binary configuration graph* (Fig. 2). The nodes of H_C represent the possible categories of the nodes of C_n and the edges the possible neighbors in C_n, considering a clockwise walk. For example, there is no edge from B_0 to R_0 because the neighbors of B_0 are both in state 1, whereas a node in R_0 is in state 0.

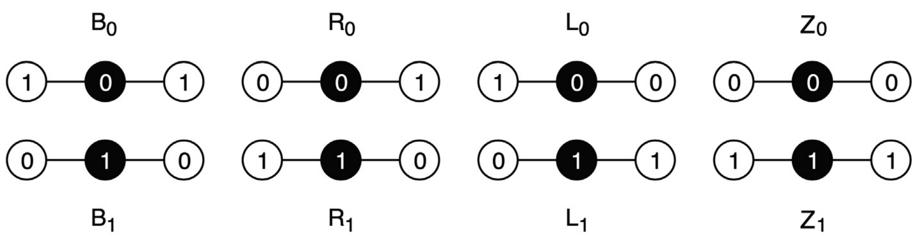

Fig. 1. Categories of a node v in C_n; node v is black and its left and right neighbors are white.

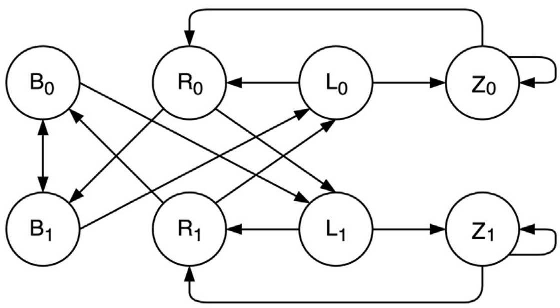

Fig. 2. The cycle binary configuration graph H_C.

Let us pick any node v in C_n and let us walk through clockwise until we return to v. Let us pick the node of H_C corresponding to the category v belongs to and follow the clockwise walk that we do on C_n also on H_C, by moving on the corresponding states. It follows that after n steps the walk on C_n will be back to v and the walk on H_C will be back to the node representing the category of v. Note that this implies that the walk on H_C is a cycle and, in general, that every cycle of length n on H_C represent a possible binary configuration of the nodes of a corresponding cycle graph C_n.

Note that every possible cycle in H_C is a combination of simple cycles (that go through each node at most once) on H_C. We prove that the structural property of the lemma holds for every simple cycle on H_C. By commutativity and associativity of addition, the property directly transfers also to composition of simple graphs. To reduce the number of simple cycles (which are 17; they are easy to find on a computer given the small size of the graph H_C, even if the problem is $\#P$-hard [4]), we avoid cycles that pass through Z_i, as $|Z_i|$ does not appear in the lemma; in fact, every cycle passing through Z_i does $L_i \to Z_i \to R_i$ and the only other outgoing edge of L_i is $L_i \to R_i$. In other words, excluding simple cycles passing through node Z_i does not have any effect on the following calculations. By also taking advantage of symmetries in i and j, all the remaining simple cycles are the following four, for which the equality of the lemma is true:

- $(B_i \to B_j)$: $|B_i|$ cancels out with $|B_j|$.
- $(B_i \to L_j \to R_j)$: $|B_i|$ cancels out with $\frac{|R_j|+|L_j|}{2}$.
- $(R_i \to L_j \to R_j \to L_i)$: $\frac{|R_i|+|L_i|}{2}$ cancels out with $\frac{|R_j|+|L_j|}{2}$.
- $(B_i \to B_j \to L_i \to R_i \to L_j \to R_j)$: $|B_i|$ cancels out with $|B_j|$; $\frac{|R_i|+|L_i|}{2}$ cancels out with $\frac{|R_j|+|L_j|}{2}$.

Theorem 4.5 (*Cycles*). *Let $G = C_n$ be the cycle on n nodes. Under the majority update rule, we have $\tau = \mathcal{O}\left(\frac{1}{\alpha} n \log n\right)$, w.h.p.*

Proof. Denote by V_i the set of nodes in state i. Given a configuration $\mathbf{x} \in \{0,1\}^n$ of C_n, let $B_i = \{v \in V_i : \forall u \in N_v, x_u \neq i\}$ and $S_i = \{v \in V_i : \exists u, w \in N_v, x_u \neq x_w\}$ (see Lemma 4.4). Let X_t be the random variable indicating the number of nodes in state 1 at time t and observe that for every k, we have:

$$\mathbf{P}(X_t = h \mid X_{t-1} = k) = \begin{cases} q_k & \text{if } h = k-1, \\ r_k & \text{if } h = k, \\ p_k & \text{if } h = k+1, \end{cases}$$

where $q_k = (1-\alpha)\left(\frac{|B_1|}{n} + \frac{1}{2}\frac{|S_1|}{n}\right)$, $p_k = \alpha\frac{n-k}{n} + (1-\alpha)\left(\frac{|B_0|}{n} + \frac{1}{2}\frac{|S_0|}{n}\right)$, and $r_k = 1 - q_k - p_k$. Therefore, the expected value of X_t, conditioned on $X_{t-1} = k$, is

$$\begin{aligned} \mathbf{E}[X_t \mid X_{t-1} = k] &= (k-1)q_k + kr_k + (k+1)p_k = k - q_k + p_k \\ &= k + \alpha\frac{n-k}{n} + \frac{1-\alpha}{n}\left(|B_0| + \frac{|S_0|}{2} - |B_1| - \frac{|S_1|}{2}\right) \\ &\overset{(a)}{=} k + \alpha\frac{n-k}{n}, \end{aligned}$$

where in derivation (a) we use Lemma 4.4. We therefore have:

$$\begin{aligned} \mathbf{E}[X_t] &= \sum_{k=0}^{n} \mathbf{E}[X_t \mid X_{t-1} = k]\mathbf{P}(X_{t-1} = k) \\ &= \sum_{k=0}^{n} \left(\alpha + \left(1 - \tfrac{\alpha}{n}\right)k\right)\mathbf{P}(X_{t-1} = k) \\ &= \alpha\sum_{k=0}^{n} \mathbf{P}(X_{t-1} = k) + \left(1 - \tfrac{\alpha}{n}\right)\sum_{k=0}^{n} k\mathbf{P}(X_{t-1} = k) \\ &= \alpha + \left(1 - \tfrac{\alpha}{n}\right)\mathbf{E}[X_{t-1}]. \end{aligned}$$

Solving this recursion with $\mathbf{E}[X_0] = 0$ we get

$$\mathbf{E}[X_t] = \alpha\sum_{i=0}^{t-1}\left(1 - \frac{\alpha}{n}\right)^i = \alpha\frac{1 - (1-\alpha/n)^t}{\alpha/n}.$$

The expected number $\mathbf{E}[n - X_t]$ of nodes in state 0 at time t is thus

$$\mathbf{E}[n - X_t] = n\left(1 - \frac{\alpha}{n}\right)^t \leqslant ne^{-\frac{\alpha}{n}t},$$

which is smaller than $\frac{1}{n}$ for $t \geqslant \frac{2}{\alpha} n \ln n$. Hence,

$$\mathbf{P}\left(\tau > \frac{2}{\alpha} n \ln n\right) = \mathbf{P}\left(n - X_{\frac{2}{\alpha} n \ln n} \geqslant 1\right) \overset{(b)}{\leqslant} \mathbf{E}\left[n - X_{\frac{2}{\alpha} n \ln n}\right] \leqslant 1/n,$$

where in (b) we use Markov's inequality.

4.3. More Low-Density Graph Families

It is not difficult to show that convergence times are also polynomial in the cases of disconnected cliques of size $\mathcal{O}(\log n)$ and trees of degree $\mathcal{O}(\log n)$. These results are summarized as the following theorem.

Theorem 4.6 (*Disconnected cliques and bounded-degree trees*). Let $G = (V, E)$ be *a set of* ℓ *disconnected cliques* $V_1, \ldots, V_\ell$ *such that* $|V_i| = \mathcal{O}(\log n)$ *for every i (respectively a tree of degree* $\mathcal{O}(\log n)$ *and such that root nodes of every subtree have at least 2 children). Then, for every constant* $\alpha \in (0, 1]$, *the expected absorption time is polynomial.*

Proof. First we start with the proof for disconnected cliques. Then, using the same notation and in part the same tools, we move to bounded-degree trees.

Disconnected cliques. Let $T_0 = 0$ and, for $j = 1, 2, 3, \ldots$, let T_j be the first round in which every node has updated its state at least once after time T_{j-1}. In the beginning of the proof of Lemma 4.2, we proved that for every j, we have that $T_j - T_{j-1} \leqslant (1+\gamma) n \ln n$ with probability at least $1 - \frac{1}{n^\gamma}$.

For a given range $[T_{j-1}, T_j]$, let $T(u)$ be the last time in which node u updated its state in the range and let X_u be the random variable such that $X_u = 1$ if u adopts opinion 1 at time $T(u)$ and 0 otherwise. Note that every X_u stochastically dominates an independent Bernoulli random variable Z_u of parameter α. Therefore, the probability for a given clique V_i of reaching absorption within the given time range is

$$\mathbf{P}(\forall u \in V_i, X_u = 1) \geqslant \mathbf{P}(\forall u \in V_i, Z_u = 1) = \alpha^{|V_i|} \geqslant \alpha^{\max_i |V_i|}.$$

Therefore, we can define a family of independent and identically distributed Bernoulli random variables $\left\{Y_i^j\right\}_{i,j}$ with parameter $p = \alpha^{\max_i |V_i|}$ and a coupling such that, for every clique V_i and every range $[T_{j-1}, T_j]$, if $Y_i^j = 1$ then V_i reached absorption within round T_j. Hence, for any given clique V_i, the checkpoint T_k in which clique V_i reaches absorption, is dominated by a geometric random variable G_i, with probability of success in each Bernoulli trial equal to $p = \alpha^{\max_i |V_i|}$ and expected value $1/p = (1/\alpha)^{\max_i |V_i|} = \mathcal{O}(\text{poly}(n))$, because $\alpha = \Omega(1)$ and $\max_i |V_i| = \mathcal{O}(\log n)$ by our hypothesis. Therefore, the expected value of the absorption time τ is upper bounded by the maximum among $\ell \leqslant n$ i.i.d. geometric random variables $\{G_i\}_i$, namely

$$\mathbf{E}[\tau] \leqslant \mathbf{E}\left[\max_{i=1,\ldots,\ell} G_i\right] \leqslant \mathbf{E}\left[\sum_{i=1}^{\ell} G_i\right] = \sum_{i=1}^{\ell} \mathbf{E}[G_i] = \frac{\ell}{p} = \mathcal{O}(\text{poly}(n)).$$

Bounded-degree trees. As a first step, we separately consider the ℓ subtrees, each consisting of a set of sibling-leaves together with their common parent. Note that, whenever such a subtree locally reaches absorption, it remains *stable* forever: leaves have their only neighbor (namely their parent) already in state 1, and the parent has the majority of their neighbors in state 1 (namely the leaves, given that the number of children is at least 2 by hypothesis). With the same reasoning used for disconnected cliques in the previous paragraph, we get that the ℓ subtrees become stable within $\mathcal{O}(\text{poly}(n))$ time steps in expectation.

Once all such structures have become stable, the ℓ parents of the subtrees will become the majority in state 1 of their respective parents. Thus, it will be sufficient to wait for the parents to update once more within the next checkpoint. Given that the same reasoning can be applied recursively for each level of the tree, from the bottom and up to the root, another $\mathcal{O}(n)$ checkpoints (i.e., as many as the maximum possible height of the tree) will suffice to globally reach absorption, still in $\mathcal{O}(\text{poly}(n))$ time steps.

5. Absorption Time for the Voter Model

As mentioned in the introduction, the voter model has received considerable attention as an opinion dynamics [31]. It may be regarded as a linearized form of the majority update rule, in the sense that, upon selection, a node selects each of the two available opinions with probability proportional to the opinion's support within the node's neighborhood. Despite such apparent similarity, the two update rules result in quite different behaviors of the biased opinion dynamics. Namely, for the voter model, absorption times to the dominant opinion are polynomial with high probability as long as $1/\alpha$ is polynomial, regardless of the underlying topology. These results are clearly at odds with those of Section 4.

The biased voter model can formally be defined as follows: Starting from some initial state $\mathbf{x}^{(0)}$, at each time t a node $u \in [n]$ is chosen uniformly at random and its opinion is updated as

$$x_u^{(t)} = \begin{cases} 1 & \text{with probability } \alpha, \\ V_G(u, \mathbf{x}) & \text{with probability } 1-\alpha, \end{cases}$$

where $\alpha \in (0,1]$ is a parameter measuring the bias toward the better opinion 1 and $V_G(u,\mathbf{x})$ is the value held in configuration $\mathbf{x}^{(t-1)} = \mathbf{x}$ by a node sampled uniformly at random from the neighborhood of node u. We assume $\mathbf{x}^{(0)} = \mathbf{0}$ for simplicity, though we remark that Theorem 5.1 below holds for any $\mathbf{x}^{(0)} \in \{0,1\}^n$.

As the proof of Theorem 5.1 highlights, the biased opinion dynamics under the voter update rule can be succinctly described by a *nonhomogeneous* Markov chain [40]. Although nontrivial to study in general, we are able to provide tight bounds in probability for the simplified setting we consider.

Theorem 5.1. Let $G = (V,E)$ be an arbitrary graph. The biased opinion dynamics with voter as update rule reaches state $\mathbf{1}$ within $\tau = \mathcal{O}(\frac{1}{\alpha} n \log n)$ steps, w.h.p.

Proof. For every node $v \in V$, the expected state of v at time $t+1$, conditioned on $\mathbf{x}^{(t)} = \mathbf{x}$ is

$$\mathbf{E}\left[x_v^{(t+1)} \mid \mathbf{x}^{(t)} = \mathbf{x}\right] = \tfrac{1}{n}\left(\alpha + \tfrac{(1-\alpha)}{d_v}\sum_{u \in N_v} x_u\right) + \left(1 - \tfrac{1}{n}\right)x_v = \tfrac{\alpha}{n} + \tfrac{1}{n}\left((1-\alpha)(P\mathbf{x})_v + (n-1)(I\mathbf{x})_v\right),$$

where $P = D^{-1}A$ is the transition matrix of the simple random walk on G (with D being the diagonal degree matrix and A the adjacency matrix of the graph) and I is the identity matrix. Considering all the nodes, we can write the vector form of the previous equation as follows:

$$\mathbf{E}\left[\mathbf{x}^{(t)} \mid \mathbf{x}^{(t-1)} = \mathbf{x}\right] = \frac{\alpha}{n}\mathbf{1} + \frac{1}{n}((1-\alpha)P + (n-1)I)\mathbf{x}.$$

This immediately implies the following equation, relating the expected states at times $t-1$ and t (with $\mathbf{E}[\mathbf{x}^{(0)}] = \mathbf{x}$):

$$\mathbf{E}\left[\mathbf{x}^{(t)}\right] = \frac{\alpha}{n}\mathbf{1} + \frac{1}{n}((1-\alpha)P + (n-1)I)\mathbf{E}\left[\mathbf{x}^{(t-1)}\right].$$

Now, consider $\mathbf{1} - \mathbf{x}^{(t)}$, the difference between the absorbing state vector $\mathbf{1}$ and the state vector at a generic time t. By the definition of $\mathbf{x}^{(t)}$, $(\mathbf{1} - \mathbf{x}^{(t)})_v \geqslant 0$ deterministically, for every v and for every t. As for the expectation of this difference, we have:

$$\begin{aligned} \mathbf{E}\left[\mathbf{1} - \mathbf{x}^{(t)}\right] &= \left(1 - \tfrac{\alpha}{n}\right)\mathbf{1} - \tfrac{1}{n}((1-\alpha)P + (n-1)I)\mathbf{E}\left[\mathbf{x}^{(t-1)}\right] \\ &= \tfrac{1}{n}((1-\alpha)\mathbf{1} + (n-1)\mathbf{1}) - \tfrac{1}{n}((1-\alpha)P + (n-1)I)\mathbf{E}\left[\mathbf{x}^{(t-1)}\right] \\ &\overset{(a)}{=} \tfrac{1}{n}((1-\alpha)P + (n-1)I)\mathbf{1} - \tfrac{1}{n}((1-\alpha)P + (n-1)I)\mathbf{E}\left[\mathbf{x}^{(t-1)}\right] \\ &= \tfrac{1}{n}((1-\alpha)P + (n-1)I)\mathbf{E}\left[\mathbf{1} - \mathbf{x}^{(t-1)}\right], \end{aligned} \tag{1}$$

where (a) follows from the fact that both matrices P and I have $\mathbf{1}$ as eigenvector with associated eigenvalue 1.

Moreover, we have

$$\frac{1}{n}((1-\alpha)P + (n-1)I) = \left(1 - \frac{\alpha}{n}\right)\widehat{P},$$

with $\widehat{P} := \frac{n-1}{n-\alpha}\left(\left(\frac{1-\alpha}{n-1}\right)P + I\right)$ being a stochastic matrix. This follows immediately by observing that both P and I are stochastic, so that all rows of $(1-\alpha)P + (n-1)I$ identically sum to $n - \alpha$. By solving the recursion in Eq. (1) we obtain

$$\mathbf{E}\left[\mathbf{1} - \mathbf{x}^{(t)}\right] = \left(1 - \frac{\alpha}{n}\right)^t \widehat{P}^t\left(\mathbf{1} - \mathbf{x}^{(0)}\right) \overset{(a)}{=} \left(1 - \frac{\alpha}{n}\right)^t \mathbf{1} - \left(1 - \frac{\alpha}{n}\right)^t \widehat{P}^t \mathbf{x}^{(0)},$$

where in (a) we use the fact that $\widehat{P}^t$ is a stochastic matrix, thus with main eigenvalue 1 and associated eigenvector $\mathbf{1}$. Next, observe that for every v, we have $\left(\widehat{P}^t \mathbf{x}^{(0)}\right)_v \geqslant 0$, so we also have $\mathbf{E}\left[1 - x_v^{(t)}\right] \leqslant \left(1 - \frac{\alpha}{n}\right)^t \leqslant e^{-\frac{\alpha}{n}t}$.

Therefore, for every time $t \geqslant \frac{2}{\alpha} n \ln n$, we have

$$\mathbf{E}\left[1 - x_v^{(t)}\right] \leqslant \frac{1}{n^2},$$

for every $v \in V$. Because the $x_v^{(t)}$'s are binary random variables, we have

$$\mathbf{P}\left(x_v^{(t)} = 0\right) = \mathbf{P}\left(1 - x_v^{(t)} = 1\right) \leqslant \mathbf{P}\left(1 - x_v^{(t)} \geqslant 1\right) \leqslant \mathbf{E}\left[1 - x_v^{(t)}\right] \leqslant \frac{1}{n^2},$$

where in the second-to-last inequality we used Markov's inequality. Hence, in $\mathcal{O}(\frac{1}{\alpha} n \log n)$ time steps the process converges to the absorbing state $\mathbf{1}$, with high probability.

Note that Theorem 5.1 implies that the convergence time is still $\mathcal{O}(n^{1+s}\log n)$ when $\alpha = \Theta(\frac{1}{n^s})$ for any $s > 0$, hence polynomial as long as s is constant.

6. Simulations

In this section, we empirically characterize the behavior of the α-biased majority dynamics in specific regimes that are not addressed by our theoretical analysis in Section 4. In particular, we have proved super-polynomial absorption times in graphs that have super-logarithmic minimum degree (Theorem 4.3) and fast absorption times in some specific graph topologies where the maximum degree is $\mathcal{O}(\log n)$. For these reasons, we are particularly interested in cases in which the average degree is $\Theta(\log n)$. To characterize more clearly how the absorption time depends on the interplay between α and the degree in this regime, in our simulations we focus on d-regular graphs.

Even if the behavior of the biased voter model is well characterized by our theoretical analysis in Section 5 (Theorem 5.1), we also conduct an empirical analysis to compare the absorption times of the α-biased voter with those of the α-biased majority while α varies. We perform a simulation by choosing a combination of a dense network topology and a small bias α, to highlight the difference in behavior between the two dynamics. In fact, as shown by Theorems 4.3 and 5.1, the absorption times would be polynomial for the former and exponential for the latter.

Experimental setup. We run the α-biased majority dynamics on regular graphs with $n = 1024$ nodes, letting d and α vary in the intervals $d \in \{2, 3, \ldots, n-1\}$ and $\alpha \in \{0.01, 0.02, \ldots, 1.0\}$, respectively. In more detail, for each pair (α, d): (1) we (almost) uniformly sample a d-regular graph G from $\mathcal{G}(n, d)$ using the algorithm by Kim and Vu [28]; (2) we perform 50 independent runs of the α-biased majority dynamics on G, for each run estimating the absorption time as explained below; (3) we take the average of the absorption times over the 50 independent runs. For each run of the α-biased majority dynamics on a graph G sampled from $\mathcal{G}(n, d)$, our estimate of the absorption time is $\min\{\tau, n^2\}$, where τ is the number of steps required for 95% of the nodes to be in state 1.[7] Note that this way, we are consistently underestimating absorption times, when the process does not converge within n^2 steps. It should be noted that, for each pair (α, d), in step (1) we only consider one sample G from $\mathcal{G}(n, d)$. This choice is consistent with our theoretical findings, in particular Theorem 4.3, suggesting that slow convergence occurs for every graph that satisfies conditions that essentially only depend on α and Δ.

Regarding the α-biased voter dynamics, instead, we only ran it in "adversarial" cases in which one might expect some deviation from theoretical predictions. In this perspective, we considered the complete graph with $n = 1024$ nodes and very small values of the bias parameter α, namely $\alpha \in \{0.0001, 0.0002, \ldots, 0.01\}$. The rationale behind these choices was highlighting the modest effect of two factors that instead significantly affect α-biased majority dynamics, namely, high network density and a small value of the bias α. In fact, the complete graph is the topology that makes the absorption time of the α-biased majority the highest; moreover, the smallest value of the bias we considered is small even with respect to network size, that is, $\alpha \approx (n\log n)^{-1}$. Similarly to what we did for biased majority, we performed 50 independent runs of the dynamics and estimated the absorption time as the average of the value $\min\{\tau, n^3\}$ measured in each run.[8]

Empirical observations. Fig. 3 provides an overview of the results of our simulations, highlighting a sharp phase transition in absorption times of the α-biased majority dynamics, with a threshold (represented by the narrow whitish line dividing the blue and red areas), which depends on α and d in a nonobvious fashion. The blue area in the heat map includes (α, d) pairs for which the dynamics rapidly converged to a consensus on opinion 1, with an average absorption time that as a function of n is compatible with $\mathcal{O}(n\log n)$, in particular, consistently smaller than $\frac{1}{\alpha}n\log n$. The red area, instead, includes pairs (α, d) for which the estimated absorption time of the dynamics was always n^2, namely configurations for which the dynamics never converged to absorption within n^2 updates; for those runs, it is possible to observe that the number of nodes in state 1 quickly reaches and remains concentrated at a value around αn. This last part is just an empirical observation (not immediately resulting from our plots), which is actually easy to motivate formally, stemming from the fact that, regardless of the current state, the node selected in any given step has a ground probability α to choose opinion 1.

The zoomed plot in the right pane of the same figure also shows that the scenario dramatically changes for $d = 2$, that is, when the graph is a cycle or a collection of disconnected cycles, as one would expect from Theorem 4.5. In this case, absorption is achieved in fewer than n^2 steps even for the smallest value of α we considered (0.01), consistently with the upper bound of Theorem 4.5.

The goal of Fig. 4 is to explore the behavior of the absorption time in the vicinity of the threshold. For this reason, the plots in this figure only refer to the combinations of (α, d) whose average absorption times were strictly less than n^2. Note that a given combination (α, d) appears in the plot even if a single run (out of the 50) reaches absorption in fewer than n^2 updates; thus, the absorption time could be dramatically underestimated. However, our goal here was to investigate how sharp the transition in absorption times is as we approach the threshold and in this respect, an underestimation only makes this transition appear smoother than it actually is.

[7] This is done to speed up the computation. In particular, given the asynchronous nature of the dynamics, when only a small fraction of nodes is still in state 0, an artificially long delay is introduced for the simple fact that nodes in state 0 are chosen less and less frequently for revising their opinions.

[8] The value n^3 instead of n^2 as in the biased majority is changed in order to handle much smaller values of α.

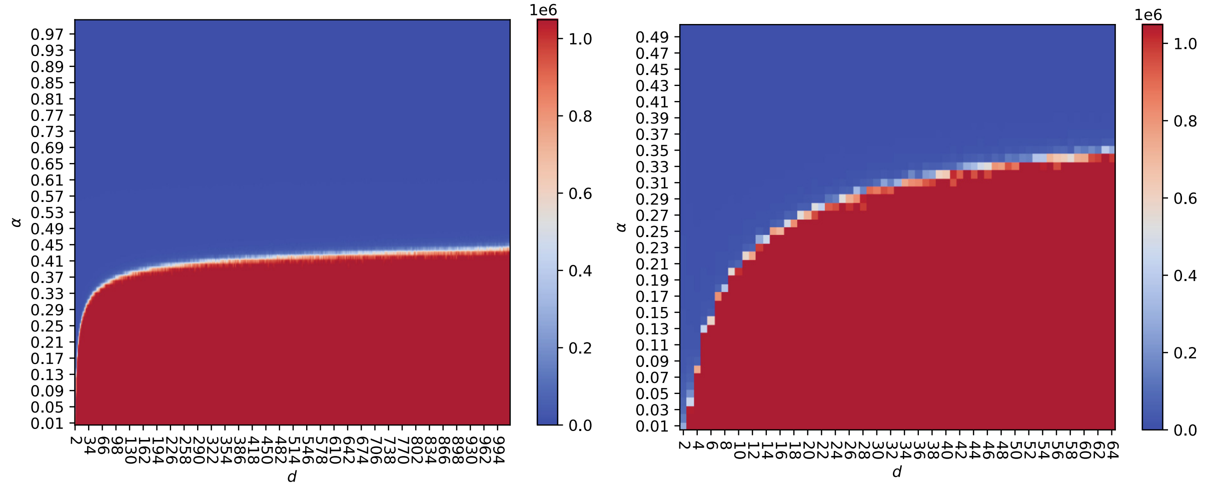

Fig. 3. Absorption time τ (color) of the biased majority dynamics on G while varying d (x-axis) and α (y-axis). On the left, the results for all values of d and α; on the right, a zoom with $d \in \{2, 3, \ldots, 2\sqrt{n}\}$ and $\alpha \in \{0.01, \ldots, 0.5\}$.

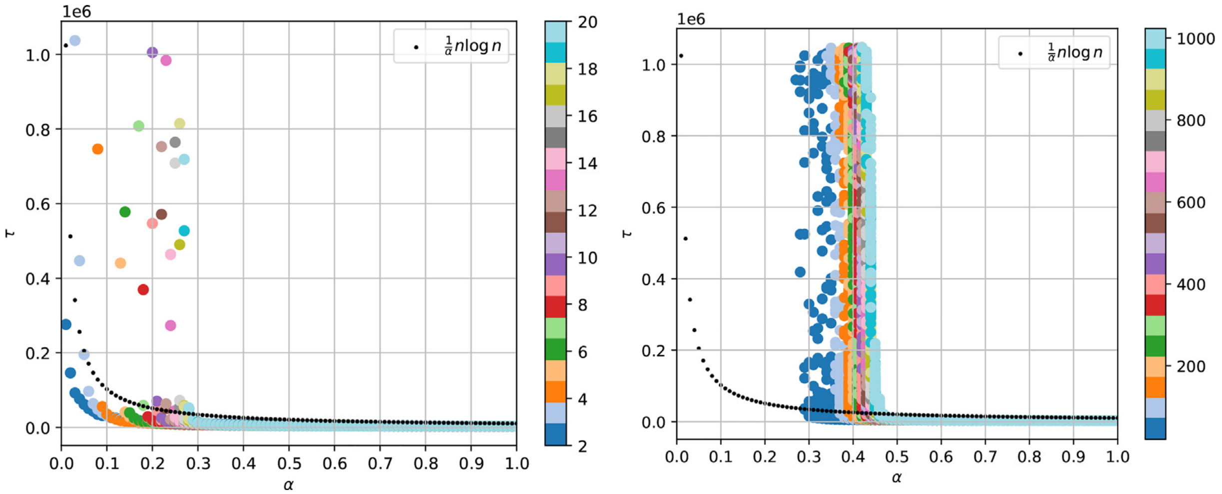

Fig. 4. Absorption time τ (y-axis) of the biased majority dynamics on G while varying α (x-axis) and d (color) vs. $\frac{1}{\alpha} n \log n$. On the left, the results for low-density graphs ($d \in \{2, 3, \ldots, 2\log_2 n\}$; on the right, the results for high-density graphs ($d \in \{1 + 2\log_2 n, \ldots, n\}$). The black dots show the value $\frac{1}{\alpha} n \log n$, which is our theoretical upper bound for the cycle.

We expect, and Fig. 4 indicates, that, for each degree d (color in the figure), there exists a threshold value for α above which the convergence is fast, whereas below it the convergence is slow; of course this value would depend on d. The figure highlights the following patterns: (1) the threshold value for α increases with the degree, which intuitively is consistent with the general insight that increased density slows the convergence of the dynamics; (2) as α approaches $1/2$ from below, fast convergence is ensured for increasingly growing values of the degree, a behavior that is ultimately consistent with Lemma 4.2 and suggesting a limiting behavior of the threshold that goes to $1/2$ for $d = n - 1$ (i.e., the complete graph) and $n \to \infty$. As commented earlier, the scenario is different in the case of $d = 2$ (leftmost blue points in the left plot) where, differently from the case of $d > 2$, the process achieved absorption for all considered values of α within n^2 steps; moreover, $d = 2$ is the only degree value for which the value of the absorption time never exceeds the black reference line $\frac{1}{\alpha} n \log n$ (i.e., our theoretical upper bound for the cycle) as α increases.

Our simulations further suggest that the lower bound we gave in Theorem 4.3 is asymptotically tight for dense graphs. In particular, recall that, whenever $\alpha = \frac{1-\varepsilon}{2}$ for a fixed constant $0 < \varepsilon < 1$, Theorem 4.3 states that if τ is the absorption time we have $\mathbf{E}[\tau] \geqslant \left(e^{\Delta\varepsilon^2/6}\right)/(6n)$, where Δ is the minimum degree. In other words, for every arbitrarily large fixed constant $K > 0$, there exists a constant $C = C(\varepsilon)$ such that $\mathbf{E}[\tau] > n^K$, whenever $\Delta > C \log n$. Put otherwise, $\mathbf{E}[\tau] > n^K$, for every fixed constant

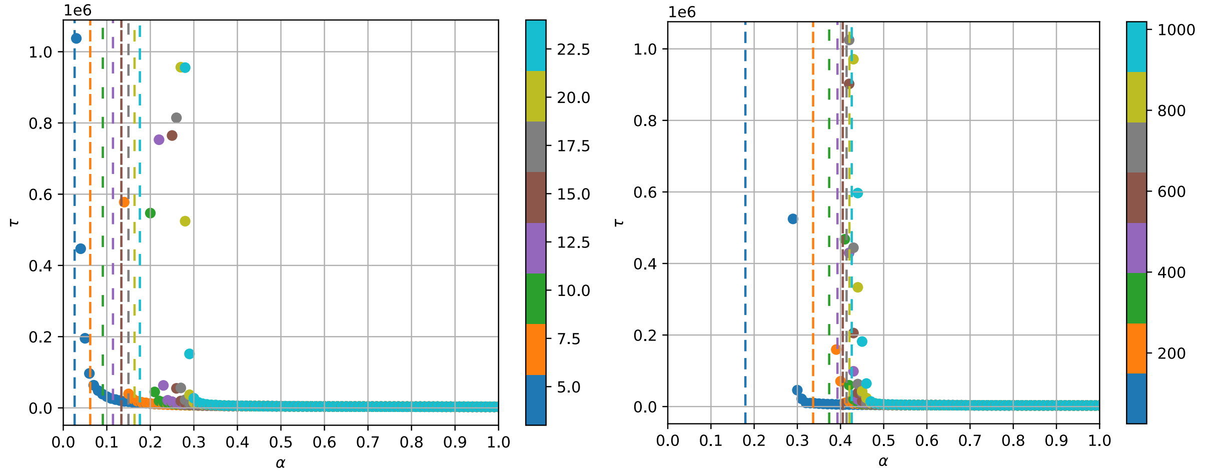

Fig. 5. Absorption time τ (y-axis) of the biased majority dynamics on G as α (x-axis) and d (color) vary. We also show as a dashed line the empirical threshold $\bar{\alpha}^*_\Delta$. In the left plot we see the results for low density graphs ($d \in \{3, \ldots, 2\log_2 n\}$, eight equally distanced values); in the right plot we see the results for high density graphs ($d \in \{2\log_2 n, \ldots, n\}$, eight equally distanced values).

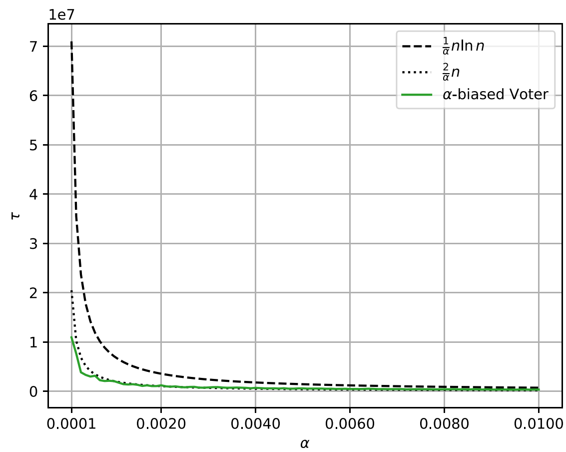

Fig. 6. Absorption time τ (y-axis) of the biased voter dynamics on the complete graph G with $n = 1024$ nodes as α (x-axis) varies. We also show as a dotted line the theoretical expected absorption time $\frac{2}{\alpha}n$ and as a dashed line the theoretical high probability upper bound on the expected absorption time $\frac{1}{\alpha}n\log n$.

$K > 0$, whenever $\alpha < \alpha^*_\Delta := \frac{1}{2}\left(1 - \sqrt{\frac{6(K+1)\ln(6n)}{\Delta}}\right)$. Although meaningless for sparse graphs,[9] this condition does apply to dense graphs in which $\Delta > 6(K+1)\ln(6n)$. For this reason we propose an empirical threshold $\bar{\alpha}^*_\Delta := \frac{1}{2}e^{-\sqrt{\frac{6(K+1)\ln(6n)}{\Delta}}}$, which is positive for all values of Δ and $\alpha^*_\Delta \leqslant \bar{\alpha}^*_\Delta \leqslant \frac{1}{2}$ (since $1 + x \leqslant e^x$ for every x).

Fig. 5 shows the empirical thresholds $\bar{\alpha}^*_\Delta$ corresponding to each value of $\Delta = d$. Even though the empirical threshold clearly differs from the actual value of α for which the process has a phase transition, especially for small values of d (left plot), the agreement becomes increasingly better as d increases (right plot). This suggests that α^*_Δ and the empirical threshold might well capture the asymptotic behavior of the phase transition in dense graphs, (asymptotically) capturing the speed of the process toward absorption. Given that $\alpha^*_\Delta \sim \bar{\alpha}^*_\Delta$ for $\Delta = \omega(\log n)$ and $n \to \infty$, these empirical observations also suggest that our lower bound provided in Theorem 4.3 might be asymptotically tight for dense graphs, although we leave this as an open question.

Fig. 6 shows the results of the simulation of the α-biased voter dynamics. The first observation is the comparison between the actual behavior of the dynamics with the theoretical prediction of Theorem 5.1. The absorption time of the α-biased voter dynamics is indeed close to its expected value (i.e., $\tau = \frac{2}{\alpha}n$); moreover, the high probability upper bound of the absorption

[9] If $\Delta < 6(K+1)\ln(6n)$ then $\alpha^*_\Delta < 0$.

time (i.e., $\tau = \frac{1}{\alpha} n \log n$), is never exceeded. The second observation regards the comparison between the actual absorption time of the α-biased voter dynamics and the lower bound on the expected absorption time of the α-biased majority dynamics (Theorem 4.3). On the one hand, the biased majority dynamics would need at least exponential time to reach absorption on a complete graph (for any $\alpha < 1/2$). On the other hand, the biased voter reaches absorption in expected time $\mathcal{O}(\frac{n}{\alpha})$ and in time $\mathcal{O}(\frac{n}{\alpha} \log n)$ with high probability.

7. Discussion and Outlook

In this paper, we considered biased opinion dynamics under two popular update rules, namely the majority [29] and the voter model [31]. Although related, these two models exhibit substantial differences in our setting. Whereas the voter model enforces a drift toward the majority opinion within a neighborhood, in the sense that this is adopted with probability proportional to the size of its support, majority is a nonlinear update rule, a feature that seems to play a crucial role in the scenario we consider. This is reflected in the absorption time of the resulting biased opinion dynamics, which is $\mathcal{O}(\frac{1}{\alpha} n \log n)$ for the voter model, regardless of the underlying topology, whereas it exhibits a far richer behavior under the majority rule, being super-polynomial (possibly exponential) in dense graphs. It may be worth mentioning that in the case of two opinions, the majority rule is actually equivalent to the (unweighted) median rule, recently proposed as a credible alternative to the weighted averaging of the DeGroot's and Friedkin-Johnsen's models [32].

A modular model. Both scenarios we studied are instantiations of a general model that is completely specified by a triple $(\mathbf{z}, \alpha, \mathcal{D})$, with $\mathbf{z}$ an initial opinion distribution, $\alpha \in (0, 1]$ a probability measuring the magnitude of the bias toward the dominant opinion, and $\mathcal{D}$ an *update rule* that specifies some underlying opinion dynamics. In more detail, a biased opinion dynamics can be succinctly described as follows.

The system starts in some state $\mathbf{x}^{(0)} = \mathbf{z}$, corresponding to the initial opinion distribution; for $t > 0$, let $\mathbf{x}^{(t-1)} = \mathbf{x}$ be the state at the end of step $t - 1$. In step t, a node v is picked uniformly at random from V and its state is updated as follows:

$$x_v^{(t)} = \begin{cases} 1 & \text{with probability } \alpha, \\ \mathcal{D}_G(v, \mathbf{x}) & \text{with probability } 1 - \alpha, \end{cases}$$

where $\mathcal{D}_G : V \times \{0, 1\}^n \to \{0, 1\}$ is the update rule.[10] When the update rule is probabilistic (as in the voter model), $\mathcal{D}_G(v, \mathbf{x})$ is a random variable, conditioned on the value $\mathbf{x}$ of the state at the end of step $t - 1$.

Remark. It is simple to see that $\mathbf{1}$ is the only absorbing state of the resulting dynamics, whenever $\alpha \neq 0$ and $\mathcal{D}$ does not allow update of an agent's opinion to one that is not held by at least one of the agent's neighbors, which is the case for many update rules in the discrete-opinion setting.

We further remark that the initial condition $\mathbf{x}^{(0)} = \mathbf{0}$ considered in this paper is not intrinsic to the model, it rather reflects scenarios (e.g., technology adoption) where a new, superior alternative to the status quo is introduced, but its adoption is possibly slowed by inertia of the system. Although the reasons behind system's inertia are not the focus of this paper, inertia itself is expressed here as a *social pressure* in the form of some update rule $\mathcal{D}_G$. Another reason for choosing a fixed initial state ($\mathbf{0}$ in our case) is being able to compare the behavior of the biased opinion dynamics under different update rules on a common basis.

Finally, it is worth mentioning that Theorem 5.1 and the upper bounds given in Section 4.3 hold regardless of the initial opinion distribution.

Outlook. This paper leaves a number of open questions. A first one concerns an accurate theoretical description of the phase-transition behavior of the dynamics under the majority update rule, as empirically observed in Section 6 on regular graphs. In fact, whereas the phenomenon is relatively well understood on dense graphs (with Lemma 4.2 and Theorem 4.3 proving a threshold phenomenon on $\alpha = 1/2$ for large n), we have not been able to provide any mathematical evidence of the same phenomenon on sparse graphs. The only exceptions are specific graph topologies for which we proved polynomial absorption time whenever α is constant, such as the cycle (Theorem 4.5), trees with degree $\mathcal{O}(\log n)$, and disconnected cliques of size $\mathcal{O}(\log n)$ (Theorem 4.6).

A further question is whether the expected absorption time is always $\mathcal{O}(n \log n)$ when $\alpha \geqslant 1/2$, irrespectively of the underlying dynamics and topology. This is clearly true for the voter model from Theorem 5.1 and it also holds for the majority model whenever the underlying network has minimum degree $\Omega(\log n)$ from Lemma 4.2.

We finally remark that our results and most results in related work apply to the case of two competing opinions. An obvious direction for further research is the extension of our results to the case of multiple opinions.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

[10] The subscript G highlights the fact that the result of the application of a given update rule $\mathcal{D}$ in general depends on both the current state and the underlying graph G. The above definition can be easily adjusted to reflect the presence of weights on the edges.

Acknowledgements

This work has been partially supported by the ERC Advanced Grant 788893 AMDROMA "Algorithmic and Mechanism Design Research in Online Markets," by the EC H2020RIA project "SoBigData++" (871042), by the MIUR PRIN project ALGADIMAR "Algorithms, Games, and Digital Markets," and by the University of "Tor Vergata" under research program "Mission: Sustainability" project ISIDE (Grant No. E81I18000110005).

References

[1] David Ahlstrom, Linda C. Wang, Groupthink and france's defeat in the 1940 campaign, Journal of Management History 15 (2) (2009) 159–177.
[2] Aris Anagnostopoulos, Luca Becchetti, Emilio Cruciani, Francesco Pasquale, and Sara Rizzo. Biased opinion dynamics: When the devil is in the details. In Christian Bessiere, editor, Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020, pages 53–59. ijcai.org, 2020..
[3] Jerry Anunrojwong, Ozan Candogan, and Nicole Immorlica. Social learning under platform influence: Extreme consensus and persistent disagreement. Available at SSRN, 2020..
[4] Sanjeev Arora, Boaz Barak, Computational Complexity - A Modern Approach, Cambridge University Press, 2009.
[5] Vincenzo Auletta, Ioannis Caragiannis, Diodato Ferraioli, Clemente Galdi, Giuseppe Persiano, Minority becomes majority in social networks, in: International conference on web and internet economics, Springer, 2015, pp. 74–88.
[6] Vincenzo Auletta, Ioannis Caragiannis, Diodato Ferraioli, Clemente Galdi, Giuseppe Persiano, Information retention in heterogeneous majority dynamics, in: International Conference on Web and Internet Economics, Springer, 2017, pp. 30–43.
[7] Vincenzo Auletta, Angelo Fanelli, Diodato Ferraioli, Consensus in opinion formation processes in fully evolving environments, AAAI Conference on Artificial Intelligence, vol. 33, 2019, pp. 6022–6029.
[8] Vincenzo Auletta, Diodato Ferraioli, and Gianluigi Greco. Reasoning about consensus when opinions diffuse through majority dynamics. In International Joint Conference on Artificial Intelligence, pages 49–55, 2018..
[9] Maryam Bahrani, Nicole Immorlica, Divyarthi Mohan, S. Matthew, Weinberg Asynchronous majority dynamics in preferential attachment trees, in: 47th International Colloquium on Automata, Languages, and Programming (ICALP 2020). Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2020.
[10] Stephen Boyd, Arpita Ghosh, Balaji Prabhakar, and Devavrat Shah. Randomized gossip algorithms. IEEE/ACM Transactions on Networking, 14(SI):2508–2530, 2006..
[11] Adrián Carro, Raúl Toral, Maxi San Miguel, The noisy voter model on complex networks, Scientific Reports 6 (1) (2016) 1–14.
[12] Peter Clifford, Aidan Sudbury, A model for spatial conflict, Biometrika 60 (3) (1973) 581–588.
[13] Adam Coates, Liangxiu Han, and Anthony Kleerekoper. A unified framework for opinion dynamics. In International Conference on Autonomous Agents and MultiAgent Systems, pages 1079–1086, 2018..
[14] Colin Cooper, Robert Elsässer, Hirotaka Ono, Tomasz Radzik, Coalescing random walks and voting on graphs, in: ACM Symposium on Principles of Distributed Computing, ACM, 2012, pp. 47–56.
[15] Colin Cooper, Robert Elsasser, Hirotaka Ono, Tomasz Radzik, Coalescing random walks and voting on connected graphs, SIAM Journal on Discrete Mathematics 27 (4) (2013) 1748–1758.
[16] Emilio Cruciani, Hlafo Alfie Mimun, Matteo Quattropani, Sara Rizzo, Phase transitions of the k-majority dynamics in a biased communication model, in: ICDCN '21: International Conference on Distributed Computing and Networking, Virtual Event, Nara, Japan, 2021, pp. 146–155.
[17] Morris H. DeGroot, Reaching a consensus, Journal of the American Statistical Association 69 (345) (1974) 118–121.
[18] Peter Donnelly, Dominic Welsh, Finite particle systems and infection models, Math. Proc. Cambridge Philos. Soc. 94 (1) (1983) 167–182.
[19] Diodato Ferraioli, Paul W Goldberg, Carmine Ventre, Decentralized dynamics for finite opinion games, Theoret. Comput. Sci. 648 (2016) 96–115.
[20] Diodato Ferraioli and Carmine Ventre. Social pressure in opinion games. In International Joint Conference on Artificial Intelligence, pages 3661–3667, 2017..
[21] Noah E Friedkin and Francesco Bullo. How truth wins in opinion dynamics along issue sequences. Proceedings of the National Academy of Sciences, 114(43), 11380–11385, 2017..
[22] Noah E Friedkin, Eugene C Johnsen, Social influence and opinions, Journal of Mathematical Sociology 15 (3–4) (1990) 193–206.
[23] James P. Gleeson, High-accuracy approximation of binary-state dynamics on networks, Phys. Rev. Lett. 107 (Aug 2011) 068701.
[24] James P Gleeson. Binary-state dynamics on complex networks: Pair approximation and beyond. Physical Review X, 3(2):021004, 2013..
[25] Yehuda Hassin, David Peleg, Distributed probabilistic polling and applications to proportionate agreement, in: International Colloquium on Automata, Languages, and Programming, Springer, 1999, pp. 402–411.
[26] Richard A Holley, Thomas M Liggett, et al, Ergodic theorems for weakly interacting infinite systems and the voter model, The Annals of Probability 3 (4) (1975) 643–663.
[27] Varun Kanade, Frederik Mallmann-Trenn, and Thomas Sauerwald. On coalescence time in graphs: When is coalescing as fast as meeting? In ACM-SIAM Symposium on Discrete Algorithms, pages 956–965, 2019..
[28] Jeong Han Kim, H.Vu. Van, Generating random regular graphs, in: Proceedings of the Thirty-Fifth Annual ACM Symposium on Theory of Computing, STOC '03, Association for Computing Machinery, New York, NY, USA, 2003.
[29] Paul L. Krapivsky, Sidney Redner, Dynamics of majority rule in two-state interacting spin systems, Phys. Rev. Lett. 90 (23) (2003) 238701.
[30] David A. Levin, Yuval Peres, Markov chains and mixing times, Am. Math. Soc. 107 (2017).
[31] Thomas Milton Liggett. Interacting particle systems, volume 276. Springer Science & Business Media, 2012..
[32] Wenjun Mei, Francesco Bullo, Ge Chen, and Florian Dörfler. Occam's razor in opinion dynamics: The weighted-median influence process. arXiv preprint arXiv:1909.06474, 2019..
[33] Mauro Mobilia, Does a single zealot affect an infinite group of voters?, Phys Rev. Lett. 91 (2) (2003) 028701.
[34] Andrea Montanari, Amin Saberi, The spread of innovations in social networks, Proc. Nat. Acad. Sci. 107 (47) (2010) 20196–20201.
[35] Stephen Morris, Contagion, Rev. Econ. Stud. 67 (2000) 57–79.
[36] A. Mukhopadhyay, R.R. Mazumdar, R. Roy, Binary opinion dynamics with biased agents and agents with different degrees of stubbornness, International Teletraffic Congress, vol. 01, 2016, pp. 261–269.
[37] Charlotte Out and Ahad N. Zehmakan. Majority vote in social networks: Make random friends or be stubborn to overpower elites. In Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21, pages 349–355. International Joint Conferences on Artificial Intelligence Organization, 8 2021..
[38] Romualdo Pastor-Satorras, Claudio Castellano, Piet Van Mieghem, Alessandro Vespignani, Epidemic processes in complex networks, Rev. Mod. Phys. 87 (Aug 2015) 925–979.
[39] Antonio F Peralta, Matteo Neri, János Kertész, and Gerardo Iñiguez. Effect of algorithmic bias and network structure on coexistence, consensus, and polarization of opinions. Physical Review E, 104(4):044312, 2021..
[40] Eugene Seneta, Non-negative matrices and Markov chains, Springer Science & Business Media, 2006.
[41] H Peyton Young. The dynamics of social innovation. Proceedings of the National Academy of Sciences, 108(Supplement 4), 21285–21291, 2011..