

# A Phase Transition for Opinion Dynamics with Competing Biases

Federico Capannoli<sup>1</sup>, Emilio Cruciani<sup>2</sup>, Hlafo Alfie Mimun<sup>3</sup>, Matteo Quattropani<sup>4</sup>

<sup>1</sup>Leiden University

<sup>2</sup>European University of Rome

<sup>3</sup>John Cabot University

<sup>4</sup>Roma Tre University

f.capannoli@math.leidenuniv.nl, emilio.cruciani@unier.it, hlafo.mimun@johncabot.edu, matteo.quattropani@uniroma3.it

## Abstract

We study the nonlinear evolution of binary opinions in a population of agents connected by a directed network, influenced by two competing forces. On the one hand agents are stubborn, i.e., have a tendency for one of the two opinions; on the other hand there is a disruptive bias that drives the agents toward the opposite opinion. The disruptive bias models external factors such as market innovations or social controllers aiming to challenge the status quo, while stubbornness reinforces the initial opinion making it harder for the external bias to drive the process toward change. Each agent updates its opinion according to a nonlinear rule that takes into account the opinions of its neighbors and the strength of the disruptive bias. We focus on random directed graphs with prescribed in- and out-degree sequences and prove that the dynamics exhibits a phase transition. When the disruptive bias is stronger than a certain critical threshold, the entire population rapidly converges to a consensus on the disruptive opinion. When the bias is weaker than this threshold, the system enters a metastable state in which only a fraction of the population adopts the new opinion, and this partial adoption persists for a long time. We explicitly characterize both the critical threshold and the long-term adoption fraction, showing that they depend only on few simple statistics of the degree sequences. Our analysis relies on a dual system of coalescing, branching, and dying particles, whose behavior is equivalent and allows a rigorous characterization of the system’s dynamics. Our results characterize the interplay between the degree of the agents, their stubbornness, and the external bias, shedding light on the tipping points of opinion dynamics in networks.

**Extended version** — <https://arxiv.org/abs/2511.09434>

## Introduction

Modeling the evolution of opinions in social and technological systems is a fundamental problem at the intersection of artificial intelligence and network science (Castellano, Fortunato, and Loreto 2009; Newman 2018). In particular, understanding how collective behaviors emerge from individual interactions is key to predicting and influencing processes such as the adoption of new technologies (Montanari and Saberi 2010) or the spread of misinformation (Del Vicario et al. 2016). Herein, we propose and analyze a nonlinear opinion dynamics model that captures the competition

between two opposing forces: an external *disruptive bias*, which pushes the system toward a new opinion, and individuals’ *stubbornness*, which weakens the effect of the external disruptive bias by reinforcing the old opinion.

The two competing biases give rise to a complex interplay between endogenous and exogenous forces in the system, which is ubiquitous in real-world systems. The disruptive bias can represent external interventions such as targeted marketing campaigns, technological innovations, or coordinated disinformation efforts (Montanari and Saberi 2010; Anagnostopoulos et al. 2022; Lesfari, Giroire, and Pérennes 2022). Stubbornness, on the other hand, reflects intrinsic properties of the population, such as entrenched cultural values, social inertia, or loyalty to existing norms or products (Yildiz et al. 2013a; Mobilia 2003). The dynamics of these competing forces provide a framework for studying tipping points, metastable states, and the long-term evolution of opinions in a wide variety of contexts.

In our model, the  $n$  nodes of a network, the agents, support a binary opinion: the *status quo*  $r$  or the *innovation*  $b$ . As discussed previously, our model takes into account two competing biases at play. First, the *disruptive bias*  $p \in [0, 1]$ , which makes each agent perceive a neighbor’s opinion as  $b$  with probability  $p$ , independently of the neighbor’s actual opinion. This parameter models external interventions or the introduction of a new technology in the network. Second, the *individuals’ stubbornness* which amplifies the probability that the agents stick to opinion  $r$ , needing to observe multiple agents with opinion  $b$  to switch. This behavior of the dynamics captures the inherent resistance or attachment of the agents to their initial state. In our dynamics, the agents sample  $s$  random neighbors and adopt opinion  $b$  only if the sampled neighbors are *all* perceived as  $b$  after accounting for the effect of the bias; otherwise, they adopt opinion  $r$ . Since the number of sampled neighbors  $s$  is greater than one, the dynamics becomes nonlinear, generating complex phenomena and phase transitions, e.g., as for the 2-Choices dynamics (Cruciani et al. 2021) which has a closely related update rule where there is no stubbornness’ effect. On the other hand, if each node would sample one single neighbor, the dynamics would be equivalent to the (biased) voter model (Holley and Liggett 1975), where each agent copies the opinion of that one randomly sampled neighbor.

While the role of network structure in shaping opinion

evolution is well acknowledged, most real-world networks are too complex and variable to allow for a clean theoretical analysis. This motivates the study of opinion dynamics on *random graphs*, which serve as idealized yet flexible models of large-scale social and technological networks. The use of randomness enables one to abstract from the microscopic details of any particular network instance and instead focus on the influence of macroscopic structural features—such as degree distributions, sparsity, or directionality—on global dynamical behavior. In this work we consider the *Directed Configuration Model* (DCM) (Van Der Hofstad 2024), that provides a natural setting to capture asymmetries in influence and information flow and to model systems where the ability to be influenced and the capacity to influence others are decoupled. This is crucial for representing many online platforms (e.g., X or Instagram), where users may have very different numbers of followers and followees. Our aim in studying biased opinion dynamics on DCM random graphs is to uncover how simple, local update rules interact with global network heterogeneity and directionality. In this context, randomness is not a bug, but a feature: it allows us to identify the key statistical features of the degree sequences that govern macroscopic phenomena, such as phase transitions in consensus dynamics, while avoiding overfitting to any specific network topology.

## Our Contribution

In many real-world social systems, individuals form opinions based on limited, potentially biased information about their peers. This is particularly evident in digital platforms, where users are exposed to content through noisy, algorithmically-curated feeds, often favoring certain narratives. Understanding how such local distortions aggregate into global consensus—or persistent disagreement—is a fundamental challenge in the study of opinion dynamics.

In this work, we introduce and analyze a novel model of *biased majority dynamics* on directed graphs. Each node in the network holds a binary opinion—*red* or *blue*—and updates asynchronously by sampling the opinions of  $s = 2$  out-neighbors; in this presentation we fix the level of stubbornness  $s$  for the sake of simplicity, but this is not restrictive as we discuss later in this section. Specifically, at Poissonian times, a node selects two out-neighbors uniformly at random and observes each as blue with probability  $p \in [0, 1]$ , independently of their true state. The node adopts opinion blue if both observed neighbors appear blue, and otherwise chooses red. This setup models a basic form of *disruptive bias*: even when red is the prevailing opinion, blue can spread if it is more visible or salient.

Despite the simplicity of the rule, the model exhibits rich and surprising behavior. While the system always eventually reaches the absorbing all-blue state when  $p > 0$ , our focus is on the timescale and nature of this convergence. In particular, we investigate how long the red opinion persists when the dynamics starts from a uniform red configuration. Our main result shows that, when the underlying graph is drawn from the DCM, the system undergoes a sharp *phase transition* as the bias parameter  $p$  varies (Theorem 3). Below a critical threshold  $p_c$ , a positive fraction of red opinions per-

sists for a long time; above it, the system rapidly reaches consensus on the blue opinion.

A key element in our analysis is a duality between the opinion dynamics and a system of labeled particles—COalescing, BRANching, and DYing (COBRAD) according to the same random environment (Proposition 1). This dual system enables us to characterize the evolution of opinions via extinction properties of particles. Crucially, we find that the macroscopic behavior of the system depends only on two simple statistics of the degree sequence: the harmonic mean of the out-degrees weighted by in-degree (denoted  $\varrho$ ), and its unweighted counterpart  $\lambda$  (see Eqs. (4) and (5)). These quantities capture how influence and exposure are distributed across the network and determine both the critical threshold  $p_c(\varrho)$  and the limiting red density  $q_*(p, \varrho, \lambda)$  in the sub-critical regime (resp. Eqs. (7) and (8)). All our theoretical statements extend to arbitrary integer levels of stubbornness  $s \geq 2$ , resulting in higher-order (in  $s$ ) degree statistics that influence the process quantitatively but not qualitatively.

Our results shed light on how structural asymmetries and local misperceptions can interact to sustain disagreement or drive consensus. Beyond the theoretical interest, this work provides a minimal framework to understand the dynamics of biased information spread in complex networks.

## Related Work

The problem we address intersects multiple areas with related settings and diverse motivations. We focus on discussing the contributions most relevant to this paper.

### Opinion Diffusion and Consensus

Opinion dynamics studies how groups of agents adjust their beliefs under the influence of peers and external factors. For a broad overview of opinion dynamics in multi-agent systems, see (Coates, Han, and Kleerekoper 2018). The classical goal is to analyze the time it takes for the agents to reach a *consensus* on some opinion. A key distinction in the field lies in how opinions are represented. The case of continuous opinions has received significant attention, particularly in social sciences and economics (DeGroot 1974; Friedkin and Johnsen 1990; Proskurnikov and Tempo 2017). In this paper, however, we focus on the case of discrete opinions.

One classical model for discrete opinions is the *majority dynamics*, where agents adopt the majority opinion among their neighbors. This rule was first studied in connection with agreement phenomena in spin systems (Krapivsky and Redner 2003). More recent studies address the role of network topology in the majority dynamics, showing under which conditions an initial majority can be overturned or not (Auletta et al. 2015; Auletta, Ferraioli, and Greco 2018; Gärtner and Zehmakan 2018).

Another classical model is the *voter model*, in which agents copy the opinion of a random neighbor. Originally inspired by studies of spatial conflicts in biology and interacting particle systems (Clifford and Sudbury 1973; Holley and Liggett 1975; Liggett 2012), it has been extensively analyzed, with results that provide tight bounds on convergence times for various graph structures (Hassin and Peleg 1999;

Cooper et al. 2013) and characterize the consensus probability on a given opinion (Donnelly and Welsh 1983).

### Stubbornness and Bias in Opinion Dynamics

A relevant area of study focuses on the impact of stubborn agents on opinion dynamics. *Stubborn* agents are generally characterized by their inclination to favor either their initial or current opinion, while *zealots* are defined as agents who permanently adhere to their initial opinion without any possibility of change. The impact of stubborn agents in the continuous-opinion dynamics has been considered for the main models (Wai, Scaglione, and Leshem 2016; Shirzadi and Zehmakan 2025), also considering external sources of bias (Out et al. 2024). Within the discrete-opinion framework, which this article examines, the influence of stubborn agents and zealots has been analyzed in the context of both the voter model (Mobilia 2003; Mobilia, Petersen, and Redner 2007; Yildiz et al. 2013b) and the majority dynamics (Auletta et al. 2017; Mukhopadhyay, Mazumdar, and Roy 2020). These studies have explored their ability to overturn an initial majority, shedding light on the conditions under which such a reversal is feasible.

Another line of work has analyzed the role of external bias toward a preferred opinion. Only a handful of analytically rigorous studies have addressed the line of research of biased opinion dynamics. These include biased adaptations of the voter model (Sood, Antal, and Redner 2008; Berenbrink et al. 2016; Mukhopadhyay, Mazumdar, and Roy 2020; Anagnostopoulos et al. 2022; Becchetti et al. 2023) and of other majority-based dynamics (Lesfari, Giroire, and Pérennes 2022; Cruciani et al. 2023; Mukhopadhyay 2024), occasionally explored in the context of evolutionary game theory (Montanari and Saberi 2010). These studies explore various models, differing in how the bias is applied to update rule or in the temporal evolution of the process, e.g., synchronous vs asynchronous updates. It is worth noting that these factors can significantly influence the behavior of the resulting dynamics. In particular, as highlighted in several recent works (Hindersin and Traulsen 2014; Cooper et al. 2018; Anagnostopoulos et al. 2022), even seemingly minor modifications to a model can lead to substantial differences in its behavior, making it difficult to predict whether the findings from one model will translate into another.

Other related works, have explored the impact of the influence exerted by small sets of *elite* nodes (Avin et al. 2017, 2019; Cruciani et al. 2021) and proposed to mitigate this influence by introducing random edges into the network topology (Out and Zehmakan 2024). Some studies, instead, analyzed the effect of uniform *noise* on opinion dynamics, rather than a bias toward a stronger/preferred opinion (d’Amore, Clementi, and Natale 2022; d’Amore and Ziccardi 2022). Finally, the role of bias, stubborn agents, and social influence has been studied in collective decision making and election scenarios (Stewart et al. 2019; Bara, Lev, and Turrini 2022; Corò et al. 2022; Faliszewski et al. 2022).

### The Role of Duality in Opinion Dynamics

Duality has proven to be a powerful analytical tool in the study of interacting particle systems and, in particular, opin-

ion dynamics on graphs. In its classical form, duality provides a correspondence between the original process and a simpler or more tractable auxiliary process evolving in a dual time direction. A prominent example is the *voter model*, which is dual to a system of *coalescing random walks* (Holley and Liggett 1975). In this dual framework, the probability that a set of agents share a common opinion at time  $t$  can be expressed as the probability that their associated random walks have coalesced by time  $t$ . This connection has been instrumental in deriving results on consensus times, fixation probabilities, and spatial correlations (Liggett 2012).

Beyond the voter model, duality has also been extended to more complex systems, including models with asymmetric interaction rules or with branching mechanisms. For example, biased voter models and majority dynamics have been analyzed using dual processes involving coalescing-branching random walks (Cooper, Radzik, and Rivera 2016). In these cases, duality provides not only computational convenience but also a probabilistic interpretation of how information or opinions propagate and interact over time. Our work contributes to the class of studies that exploit duality to analyze opinion dynamics, by relating the evolution of opinions to a tractable dual process that captures the flow of influence across the underlying graph, in a setting characterized by competing biases.

### Model and Results

Let  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$  be a *directed graph* with  $|\mathcal{V}| = n$ , and in which every vertex has out-degree at least 1. For a given vertex  $x \in \mathcal{V}$ , we call  $\mathcal{N}_x = \{y \in \mathcal{V} \mid (x, y) \in \mathcal{E}\}$  the set of out-neighbors of  $x$ . Each vertex is equipped with a binary opinion which we will identify as blue,  $b$ , or red,  $r$ . For  $x \in \mathcal{V}$  and  $t \geq 0$  we write  $c_x(t) \in \{b, r\}$  for the opinion of vertex  $x$  at time  $t$ . We attach to each vertex an independent Poisson process of rate 1. For a given vertex  $x \in \mathcal{V}$ , at each arrival of the corresponding Poisson process the vertex selects  $s = 2$  (the *stubbornness*) out-neighbors uniformly at random, with replacement. The vertex then observes the color of such two neighbors: regardless of their actual colors, there is a probability  $p \in [0, 1]$  (the *disruptive bias*) that the observed neighbor is seen blue. Each observation is independent of the others. After the observations the vertex assumes the color blue only if *both* neighbors are seen blue, otherwise it assumes color red. See an example in Figure 1.

Notice that the presence of the bias  $p$  creates a *frustration* in the system: on the one hand, each vertex has a tendency on




Figure 1: Graphical example of the update rule: vertex  $u$  samples two neighbors  $v, w$ ; the bias is applied on  $v$ , that is seen as *blue*, but not on  $w$ . Vertex  $u$  updates to blue as it sees blue twice (effect of bias on  $v$  and  $w$  is blue).

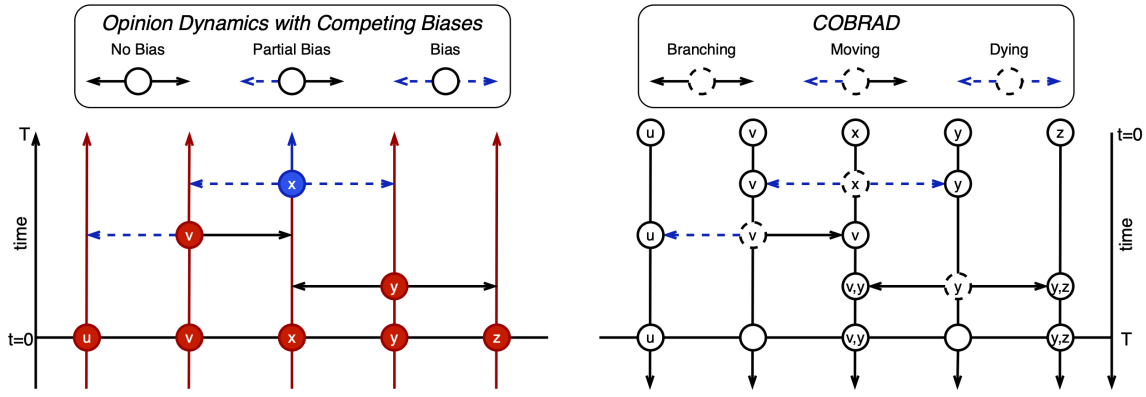


Figure 2: Graphical representation of the opinion dynamics with competing biases (left) and of COBRAD (right). On the left, in chronological order,  $y$  samples  $x, z$  without the effect of the bias,  $v$  samples  $u, x$  with the effect of the bias on  $u$ , and  $x$  samples  $v, y$  with the effect of the bias on both. On the right, in chronologically reverse order, the particle labeled  $x$  dies, the particle labeled  $v$  moves to vertex  $x$ , and the particle labeled  $y$  branches to both vertices  $x$  and  $z$ . At time 0 in COBRAD there are no particles labeled  $x$ , hence vertex  $x$  in the opinion dynamics with competing biases at time  $T$  has color  $b$ .

adopting the opinion red, i.e., is *stubborn*, since ties are broken in favor of this opinion; on the other hand, the presence of the *bias* favors the opinion blue. It is clear that, regardless of the initial configuration, the system will be eventually absorbed in a configuration in which every vertex is blue. In what follows we will denote by  $\xi_t \in \{r, b\}^{\mathcal{V}}$  the opinion configuration at time  $t \geq 0$ , and use the shorthand notation  $\xi^r$  (resp.  $\xi^b$ ) to denote the configuration in which every vertex has opinion  $r$  (resp.  $b$ ). We will further write  $R_t$  to denote the set of vertices sharing the opinion  $r$  at time  $t \geq 0$ , and we will assume that the initial configuration is  $\xi^r$ . We are interested in understanding the time it takes for the system to reach the configuration  $\xi^b$  as a function of the bias,  $p$ , and on the geometry of the underlying graph,  $\mathcal{G}$ .

It will be convenient to describe our opinion dynamics by means of a so-called *graphical construction*. To each  $x \in \mathcal{V}$  we associate a marked Poisson process of rate 1. For each  $x \in \mathcal{V}$  and for each arrival  $t \geq 0$  of the associated Poisson process we attach a mark

$$(N_1(x, t), N_2(x, t), M_1(x, t), M_2(x, t)) \in \mathcal{N}_x \times \mathcal{N}_x \times \{0, 1\} \times \{0, 1\}. \quad (1)$$

More precisely, the four components of the marks are sampled jointly independently, with the first two being marginally uniformly distributed and the last two being distributed as  $\text{Bern}(p)$ .  $N_1(x, t)$  and  $N_2(x, t)$  correspond to the neighbors of  $x$  which are observed at time  $t$ , while the other two marks represent a bias in that observation: if  $M_1(x, t) = 1$  (resp.  $M_2(x, t) = 1$ ), then the neighbor  $N_1(x, t)$  (resp.  $N_2(x, t)$ ) is seen blue regardless of its current color. Then, the vertex  $x$  becomes blue if both the sampled neighbors are seen blue. Otherwise, if  $x$  sees at least a red neighbor, then it will become red.

### The Dual Particle System

A key observation is that the opinion dynamics described above is dual to a system of COalescing, BRANching and

Dying labeled particles, which we will denote by COBRAD. At time 0 each particle sits on a vertex, and it is labeled by the label of such a vertex. Moreover, when two particles with the same label sit on the same vertex, we will treat them as a single particle. Therefore, the process  $(\eta_t)_{t \geq 0}$  has state space  $\mathcal{V}^{\mathcal{V}}$ . In words, for every  $x \in \mathcal{V}$ ,  $\eta_t(x) \subset \mathcal{V}$  denotes the set of the labels of the particles sitting at  $x$  at time  $t$ . Moreover, for any  $y \in \mathcal{V}$  we will write  $y \in \eta_t$  to denote that there exists some  $x \in \mathcal{V}$  such that  $y \in \eta_t(x)$ . The process evolves by means of the same marked Poisson processes used for the opinion dynamics. In particular, at the arrivals of the Poisson process associated to vertex  $x$ , one of the following scenarios occur:

- If  $M_1(x, t) = M_2(x, t) = 0$ , then each particle at  $x$  undergoes a branching, i.e., it is removed from  $x$  and it produces two particles (with the same label) which are placed at  $N_1(x, t)$  and at  $N_2(x, t)$ .
- If  $M_1(x, t) = 0$  while  $M_2(x, t) = 1$ , then each particle at  $x$  moves to  $N_1(x, t)$ .
- Similarly, if  $M_2(x, t) = 0$  while  $M_1(x, t) = 1$ , then each particle at  $x$  moves to  $N_2(x, t)$ .
- Finally, if  $M_1(x, t) = M_2(x, t) = 1$  all the particles sitting at  $x$  die.

Notice that, regardless of the initial configuration, the particles system will undergo an extinction in finite time *almost surely*, i.e., with probability 1.

For a given directed graph, and for a prescribed time horizon  $T > 0$  we sample the marked Poisson processes that will be used to construct the opinion dynamics up to time  $T$ . Then, we consider the COBRAD particle system which uses the same marked Poisson process, but backward in time. More precisely, for some  $A \subseteq \mathcal{V}$  we start with the configuration  $\eta_0 = \eta^A$  having at site  $x \in \mathcal{V}$  a single particle (with label  $x$ ) if and only if  $x \in A$ , and otherwise  $x$  is empty at

time 0. We consider the marked Poisson process

$$(\widehat{\text{PP}}_t(x))_{x \in \mathcal{V}, t \in [0, T]} = (\text{PP}_{T-t}(x))_{x \in \mathcal{V}, t \in [0, T]}.$$

The following fact is an immediate consequence of the latter construction.

**Proposition 1.** *For any  $A \subseteq V$ ,  $C \subseteq A$  and  $T > 0$*

$$\begin{aligned} \mathbf{P} \left( \bigcap_{x \in C} \{c_x(t) = b\} \cap_{x \in A \setminus C} \{c_x(t) = r\} \mid \xi_0 = \xi^r \right) \\ = \mathbf{P} \left( \bigcap_{x \in C} \{x \notin \eta_t\} \cap_{x \in A \setminus C} \{x \in \eta_t\} \mid \eta_0 = \eta^A \right). \end{aligned} \quad (2)$$

In other words, a given subset of vertices  $C \subseteq A$  has opinion  $b$  at time  $t$  if and only if in the associated particle systems the labels in  $C$  are extinct at time  $t$ . See Figure 2 for a visual representation.

### The Directed Configuration Model (DCM)

Despite the fact that the opinion dynamics described above and its dual particle system are well defined on any directed (or undirected) graph, in this work we focus on a benchmark model known as *Directed Configuration Model* (Van Der Hofstad 2024). Similarly to its undirected counterpart, (Bollobás 1980; Van Der Hofstad 2024), such a random digraph model is defined as follows: for every fixed graph size,  $|\mathcal{V}| = n$ , the in- and out-degree sequences are prescribed as parameters of the model,  $(d_x^-, d_x^+)_{x \in \mathcal{V}}$ . We will be interested in the large graph limit,  $n \rightarrow \infty$ , so the asymptotic notation is referring to such a limit and the explicit dependence on  $n$  will be dropped from the notation. We will further assume that such degree sequences satisfy the following requirements

$$\begin{cases} \sum_{x \in \mathcal{V}} d_x^- = \sum_{x \in \mathcal{V}} d_x^+ =: m; \\ \min_{x \in \mathcal{V}} d_x^+ \geq 2; \\ \Delta := \max_{x \in \mathcal{V}} d_x^+ = O(1); \\ \frac{1}{n} \sum_{x \in \mathcal{V}} (d_x^-)^{2+\delta} = O(1) \quad \text{for some } \delta > 0. \end{cases} \quad (3)$$

Before commenting on these requirements, let us explain how to obtain a random graph out of the choice of such parameters: we attach to each vertex,  $x \in \mathcal{V}$ ,  $d_x^-$  labeled tails and  $d_x^+$  labeled heads and then take a uniformly random bijection from the set of tails to the set of heads. The first requirement is needed to guarantee that such a random bijection is well defined and it can be projected into a (multi-)digraph<sup>1</sup>. The second requirement guarantees that, with high probability, the digraph obtained by this procedure admits a unique strongly-connected component, see (Bordenave, Caputo, and Salez 2018; Cai and Perarnau 2023; Cai et al. 2023). The last two requirements are of technical nature; nevertheless they are quite reasonable for a model of a real-world social network: thinking, e.g., at a social network like *Instagram*, the number of people you follow is somehow constrained to be a bounded quantity, while it is natural to assume that there are few *influencers* who have an enormous number of *followers*. In a sense, the DCM serves as a benchmark for digraphs with given degree sequences,

<sup>1</sup>We will in fact assume that multiple edges with the same source-destination are allowed, as well as self-loops. Despite that, we will omit the prefix *multi*-.

since it can be thought of as the *maximum entropy ensemble* for graphs with such in- and out-degrees. Despite the size,  $2n$ , of the model parameter, it is known that the behavior of several diffusion models and opinion dynamics on these random graphs actually depends only on a few simple statistics of the degree sequences, see (Bordenave, Caputo, and Salez 2018; Avena et al. 2024; Caputo and Quattropiani 2020). In the setup of this work, we aim at understanding which are the features of the degree sequence that have an impact on the behavior of our biased opinion dynamics. Perhaps surprisingly, we will show that the behavior is in fact dictated only by the following statistics:

$$\varrho_n := \sum_{x \in \mathcal{V}} \frac{d_x^-}{m} \frac{1}{d_x^+}, \quad \lambda_n := \frac{1}{n} \sum_{x \in \mathcal{V}} \frac{1}{d_x^+}, \quad (4)$$

namely,  $\varrho$  is the average of the inverse out-degree of a vertex sampled with probability proportional to its in-degree; while  $\lambda$  is the average of the same quantity but with respect to the uniform distribution. Notice that, thanks to (3) we have  $\varrho, \lambda \in (0, \frac{1}{2}]$ . It is worth to point out that the key role of the statistics  $\varrho$  already appeared in other works investigating stochastic processes on the DCM (Bordenave, Caputo, and Salez 2018; Avena et al. 2024), and has also been related to the spectral properties of such random graph (Coste 2021). Clearly, the quantities in (4) depend on  $n$ , therefore, in order to avoid inconsistencies, we will assume that the (sequence of) degree sequences is such that

$$\lim_{n \rightarrow \infty} \varrho_n = \varrho, \quad \lim_{n \rightarrow \infty} \lambda_n = \lambda. \quad (5)$$

### Results

Our main result is the identification of a phase transition on the behavior of the biased opinion dynamics described above when taking place on the DCM random graph. To phrase our theorem let us start with a definition.

**Definition 2.** *We will call slowly diverging sequence a positive sequence  $(f_n)_{n \in \mathbb{N}}$  such that*

$$\liminf_{n \rightarrow \infty} f_n = +\infty, \quad \limsup_{n \rightarrow \infty} \frac{f_n}{\log(n)} = 0. \quad (6)$$

**Theorem 3.** *Consider, for all  $n \in \mathbb{N}$ , the degree sequences  $(d_x^-, d_x^+)_{x \in \mathcal{V}}$  satisfying (3) and (5) and call*

$$p_c(\varrho) := \frac{\sqrt{1-\varrho} - (1-\varrho)}{\varrho}, \quad (7)$$

$$q_*(p, \varrho, \lambda) := \left(1 - \frac{p^2}{(1-p)^2(1-\varrho)}\right) \left(1 - \frac{p^2(\lambda-\varrho)}{1-\varrho}\right). \quad (8)$$

Take  $(t_n)_{n \in \mathbb{N}}$  an arbitrary slowly diverging sequence and call  $\mathbf{P}$  the joint law of the DCM and the opinion dynamics. Then:

- if  $p \geq p_c$  then for all  $\varepsilon > 0$

$$\lim_{n \rightarrow \infty} \mathbf{P} \left( \left| \frac{R_{t_n}}{n} \right| < \varepsilon \mid \xi_0 = \xi^r \right) = 1;$$

- if  $p < p_c$  then for all  $\varepsilon > 0$

$$\lim_{n \rightarrow \infty} \mathbf{P} \left( \left| \frac{R_{t_n}}{n} - q_* \right| < \varepsilon \mid \xi_0 = \xi^r \right) = 1.$$

**Remark 4.** *It is crucial to note that  $q_*(p, \varrho, \lambda) \in [0, 1]$  for all  $p < p_c(\varrho)$  and  $\lambda, \varrho \in (0, 1/2]$ , and that the function  $p \mapsto q_*(p, \varrho, \lambda)$  is decreasing regardless of  $\lambda, \varrho \in (0, 1/2]$ .*

## Interpretation of the Results

Our main theorem reveals that the behavior of the biased opinion dynamics undergoes a *sharp phase transition* as a function of the bias parameter  $p$ , with a critical threshold  $p_c(\varrho)$  that depends on the structure of the underlying directed graph. Below this threshold, a persistent density of red opinions survives for a long time; above it, the system rapidly converges to full consensus on blue—even if all individuals initially hold the red opinion.

This result offers a number of conceptual insights. First, it highlights the *nontrivial interplay between bias and network structure*. The fact that the critical bias depends on the structural parameter  $\varrho$ —which reflects how in-degrees correlate with out-degrees—indicates that the same level of local misperception may have dramatically different global effects depending on how influence and exposure are distributed across the network. In particular, the system is more susceptible to blue consensus (i.e., has lower  $p_c$ ) when individuals with many incoming connections tend to have fewer outgoing ones (i.e., has larger  $\varrho$ ), cf. Figure 3.

Second, the limiting fraction of red individuals in the subcritical regime is given explicitly by the quantity  $q_*(p, \varrho, \lambda)$ , showing that *partial consensus (or polarization)* is not a transient phenomenon but a stable long-time outcome under certain conditions. In other words, *competing biases alone are sufficient to sustain diversity of opinions*.

Third, and perhaps most surprisingly, the critical threshold and the limiting behavior of the system depend only on *aggregated statistics*  $\varrho$  and  $\lambda$ , rather than on the full degree sequence. This universality suggests that the *global behavior is insensitive to microscopic details*, and instead dictated by a coarse balance between influence and susceptibility. From a modeling perspective, this greatly simplifies the understanding of consensus dynamics on large directed graphs: if one is convinced that the structure of the network in analysis is dictated only by its first-order properties (i.e., the degrees), then one can reason effectively in terms of just two scalar quantities. This can be observed in Figure 4.

Finally, from a broader viewpoint, our findings emphasize how *individual-level misperceptions* (i.e., biased observations) can amplify and propagate across the network in a highly nonlinear way. Even a small probability of misinterpreting neighbors' opinions can lead to a tipping point where the majority flips. This provides a potential explanation for the emergence of dominant narratives or widespread misinformation in digital platforms, where the structure of influence is asymmetric and noisy observations are the norm.




Figure 3: On the  $x$ -axis have the value of  $\varrho$ , while on the  $y$ -axis the value of  $p$ . The blue (resp. red) area represents couples  $(\varrho, p)$  for which we have  $p > p_c(\varrho)$  (resp.  $p < p_c(\varrho)$ ).




Figure 4: We run the biased opinion dynamics on three DCM graphs with  $n = 10^4$  vertices: in the first all vertices have in- and out-degrees equal to 6 (red); in the second, half of the vertices have in-degree 10 and out-degree 5, and the other way around for the other half (green); in the third, half of the vertices have in-degree 10 and out-degree 2, and vice versa for the second half (blue). Notice that the red and the green graph share the same value of  $\varrho = \frac{1}{6}$  (hence, they have the same  $p_c \approx 0.477$ ) while the red and the blue graph have the same number of edges (but in the blue case  $\rho = \frac{13}{30} > \frac{1}{6}$ ). Notice also that the value of  $\lambda$  is  $\frac{1}{6}$ ,  $\frac{3}{20}$  and  $\frac{3}{10}$ , in the blue, red and green case, respectively. The value of the critical threshold in the case of the blue graph is  $p_c \approx 0.429$ . On the left a simulation for  $p = 0.3$ , which is subcritical in all the three cases; on the right a simulation for  $p = 0.45$ , which is subcritical for the red and the green graphs, but supercritical for the blue one. The dashed lines indicate the theoretical prediction of the long-term density of red opinions, i.e.,  $q_*(p, \varrho, \lambda)$ . Notice that in the case  $p = 0.3$  (left) we have  $q_*(\frac{3}{10}, \frac{1}{6}, \frac{1}{6}) \approx 0.7796$  (red),  $q_*(\frac{3}{10}, \frac{1}{6}, \frac{3}{20}) \approx 0.7810$  (green) and  $q_*(\frac{3}{10}, \frac{1}{6}, \frac{3}{10}) \approx 0.6902$  (blue); while in the case  $p = 0.45$  (right) we have  $q_*(\frac{9}{20}, \frac{1}{6}, \frac{1}{6}) \approx 0.1967$  (red) and  $q_*(\frac{9}{20}, \frac{1}{6}, \frac{3}{20}) \approx 0.1975$  (green).

## The Associated Random Tree

At the core of our analysis stands the following intuition: it is well-known that sparse random graphs can be locally approximated by random trees, whose law depends on the underlying degree sequence. For the precise mathematical formulation of this fact we refer to (Van Der Hofstad 2024) and, more specifically to (Bordenave, Caputo, and Salez 2018) for the case of the DCM. In this section we focus on the analysis of the COBRAD process on such approximating random trees, showing how the key quantities  $p_c$  and  $q_*(p)$  appear in the analysis. In the supplemental material, which is omitted for the sake of space, we show how to translate the analysis in this section on a result for the random graph model. Therefore, we do not aim at proving here in full details the result in Theorem 3, but rather on providing an intuition for its validity.

The key idea can be summarized as follows. Let us focus on a given vertex  $x \in \mathcal{V}$ . Thanks to the duality with the COBRAD process we know that the probability that  $c_x(t) = r$  coincides with the probability that the COBRAD process initialized with a single particle at  $x$  is not extinct at time  $t$ . Now, if we approximate the local out-neighborhood of  $x$  by a random tree, we can approximate the probability of survival of the COBRAD process at time  $t$  on the graph with the survival of the same process on such an approximating random tree.

Fix  $n \geq 1$  and consider the in- and out-degree sequences  $(d_x^-, d_x^+)_{x \in \mathcal{V}}$ . We construct the approximating random directed tree as follows: the root has a deterministic number of children,  $d_{\text{root}}$ , while each other vertex on the tree has a number of children,  $X$ , such that the probability that  $X = k$  corresponds to the probability that, sampling a vertex in  $\mathcal{V}$  with probability proportional to its in-degree, we obtain a vertex with out-degree  $k$ . In other words, called  $P_n$  the law of the offspring distribution, we have

$$P_n(X = k) = \sum_{x \in \mathcal{V}} \frac{d_x^-}{m} \mathbb{1}_{\{d_x^+ = k\}}, \quad k \in [2, \Delta]. \quad (9)$$

Once the tree has been constructed, we consider the COBRAD process starting with a single particle at the root. Being the tree directed, the particles can only perform downward transitions along the tree. In particular, if a particle ever visits a certain vertex of the tree having degree  $d$ , then such a particle will eventually:

- Die with probability

$$p_0(d) = p^2;$$

- Move to one of the children with probability

$$p_1(d) = (1-p) \left[ 2 \left( 1 - \frac{1}{d} \right) p + \frac{1}{d} (1+p) \right];$$

- Branch into two new particles going on two different children of the current vertex with probability

$$p_2(d) = \left( 1 - \frac{1}{d} \right) (1-p)^2.$$

This process can therefore be seen as a *percolation* on the random tree, where each edge is retained if and only if a particle eventually crosses the edge. Called  $\mathbb{P}_n$  the joint law of the tree construction and of the COBRAD, we can reinterpret the percolated tree as a random subtree in which each vertex can only have 0, 1 or 2 children. For the root, the probability of each of these outcomes is given by  $p_0(d_{\text{root}})$ ,  $p_1(d_{\text{root}})$  and  $p_2(d_{\text{root}})$ , while for all the other vertices such probabilities are given by

$$\begin{aligned} p_0 &= \sum_{x \in \mathcal{V}} \frac{d_x^-}{m} p_0(d_x^+) = p^2, \\ p_1 &= \sum_{x \in \mathcal{V}} \frac{d_x^-}{m} p_1(d_x^+) = (1-p)(2p - p\varrho_n + \varrho_n), \\ p_2 &= \sum_{x \in [n]} \frac{d_x^-}{m} p_2(d_x^+) = (1-\varrho_n)(1-p)^2. \end{aligned}$$

We are now interested in evaluating the probability that such a percolated tree is infinite. Given the special role of the root, we start by looking at the probability that the percolated tree starting at one of the root's children is infinite (so that every vertex on this subtree has the same offspring distribution). By the classical theory of random trees (see, e.g., (Lyons and Peres 2017)) one has that the probability that such tree

is infinite is zero as soon as the expected number of children is smaller than or equal to 1, i.e.,

$$p_1 + 2p_2 \leq 1 \iff p \geq p_c(\varrho_n). \quad (10)$$

More generally, the probability that such a tree is finite can be obtained by looking at the smallest solution of the equation

$$p_0 + p_1 z + p_2 z^2 = z, \quad (11)$$

that is

$$z_*(p, \varrho_n) = \frac{p^2}{(1-p)^2(1-\varrho_n)} \wedge 1. \quad (12)$$

Note that the above minimum is 1 if and only if  $p \geq p_c(\varrho_n)$ .

Keeping into account the role of the root, we have that the percolated tree starting at the root is finite with probability

$$\hat{z}_*(d_{\text{root}}, p, \varrho_n) = p_0(d_{\text{root}}) + p_1(d_{\text{root}})z_*(p, \varrho_n) + p_2(d_{\text{root}})z_*(p, \varrho_n)^2. \quad (13)$$

At this point, using such an approximation, the density of red vertices at time  $t$  is given by adding up the probabilities in (13) when the root is each  $x \in \mathcal{V}$ . As a consequence:

$$\frac{|R_t|}{n} \approx \frac{1}{n} \sum_{x \in \mathcal{V}} (1 - \hat{z}_*(d_x^+, p, \varrho_n)) = q_*(p, \varrho_n, \lambda_n), \quad (14)$$

where the latter quantity is defined as in (8).

## Conclusion and Future Work

We have introduced a simple yet expressive model of opinion dynamics with competing biases and shown that exhibits a sharp phase transition driven by the strength of the disruptive bias. Our analysis on the Directed Configuration Model reveals that long-term disagreement (or rapid consensus) can emerge purely from the interplay between local misperceptions and global network structure.

The duality with a particle system allowed us to characterize the metastable behavior of the process and to identify the critical threshold  $p_c(\varrho)$  and limiting red density  $q_*(p, \varrho, \lambda)$  in the subcritical regime. Remarkably, both quantities depend only on coarse statistics of the degree sequence, suggesting a form of universality in the large-network limit.

These findings highlight the power of minimal stochastic models in capturing essential features of social dynamics under noise and structural asymmetry.

On the mathematical side, an intriguing open problem is to understand the precise time to consensus in the large-graph limit, particularly in relation to metastability and sharp transition windows. In the spirit of the mathematical literature on the contact process on random graphs (Mourrat and Valesin 2016; Mountford et al. 2016), we formulate the following conjecture:

**Conjecture 5.** *In the same setup as Theorem 3*

- if  $p \geq p_c$  then there exists some  $\bar{c}(p) > 0$  such that

$$\lim_{n \rightarrow \infty} \mathbf{P}\left(R_{\bar{c} \log(n)} = \emptyset \mid \xi_0 = \xi^r\right) = 1;$$

- if  $p < p_c$  then there exists some  $\underline{c}(p) > 0$  such that

$$\lim_{n \rightarrow \infty} \mathbf{P}\left(R_{e \underline{c} n} \neq \emptyset \mid \xi_0 = \xi^r\right) = 1.$$

Moreover, called  $\tau_b = \inf\{t \geq 0 \mid \xi_t = \xi^b\}$ ,

$$\lim_{n \rightarrow \infty} \mathbf{P}\left(\tau_b > t \cdot \mathbf{E}[\tau_b] \mid \xi_0 = \xi^r\right) = e^{-t}.$$

## References

- Anagnostopoulos, A.; Becchetti, L.; Cruciani, E.; Pasquale, F.; and Rizzo, S. 2022. Biased opinion dynamics: when the devil is in the details. *Information Sciences*, 593: 49–63.
- Auletta, V.; Caragiannis, I.; Ferraioli, D.; Galdi, C.; and Persiano, G. 2015. Minority becomes majority in social networks. In *International conference on web and internet economics*, 74–88. Springer.
- Auletta, V.; Caragiannis, I.; Ferraioli, D.; Galdi, C.; and Persiano, G. 2017. Information retention in heterogeneous majority dynamics. In *International Conference on Web and Internet Economics*, 30–43. Springer.
- Auletta, V.; Ferraioli, D.; and Greco, G. 2018. Reasoning about Consensus when Opinions Diffuse through Majority Dynamics. In *International Joint Conference on Artificial Intelligence*, 49–55.
- Avena, L.; Capannoli, F.; Hazra, R. S.; and Quattropani, M. 2024. Meeting, coalescence and consensus time on random directed graphs. *Annals of Applied Probability*, 34(5): 4940–4997.
- Avin, C.; Lotker, Z.; Mizrahi, A.; and Peleg, D. 2019. Majority vote and monopolies in social networks. In *Proceedings of the 20th International Conference on Distributed Computing and Networking*, 342–351. New York, NY, USA: Association for Computing Machinery.
- Avin, C.; Lotker, Z.; Peleg, D.; Pignolet, Y.-A.; and Turkel, I. 2017. Elites in Social Networks: An Axiomatic Approach. In Shmueli, E.; Barzel, B.; and Puzis, R., eds., *3rd International Winter School and Conference on Network Science*, 75–87. Cham: Springer.
- Bara, J.; Lev, O.; and Turrini, P. 2022. Predicting voting outcomes in the presence of communities, echo chambers and multiple parties. *Artificial Intelligence*, 312: 103773.
- Becchetti, L.; Bonifaci, V.; Cruciani, E.; and Pasquale, F. 2023. On a Voter Model with Context-Dependent Opinion Adoption. In *32nd International Joint Conference on Artificial Intelligence, IJCAI*, 38–45. ijcai.org.
- Berenbrink, P.; Giakkoupis, G.; Kermarrec, A.; and Mallmann-Trenn, F. 2016. Bounds on the Voter Model in Dynamic Networks. In *43rd International Colloquium on Automata, Languages, and Programming*, 146:1–146:15.
- Bollobás, B. 1980. A probabilistic proof of an asymptotic formula for the number of labelled regular graphs. *European Journal of Combinatorics*, 1(4): 311–316.
- Bordenave, C.; Caputo, P.; and Salez, J. 2018. Random walk on sparse random digraphs. *Probab. Theory Related Fields*, 170(3-4): 933–960.
- Cai, X. S.; Caputo, P.; Perarnau, G.; and Quattropani, M. 2023. Rankings in directed configuration models with heavy tailed in-degrees. *Ann. Appl. Probab.*, 33(6B): 5613–5667.
- Cai, X. S.; and Perarnau, G. 2023. The diameter of the directed configuration model. *Ann. Inst. Henri Poincaré Probab. Stat.*, 59(1): 244–270.
- Caputo, P.; and Quattropani, M. 2020. Stationary distribution and cover time of sparse directed configuration models. *Probability Theory and Related Fields*, 178: 1011–1066.
- Castellano, C.; Fortunato, S.; and Loreto, V. 2009. Statistical physics of social dynamics. *Reviews of modern physics*, 81(2): 591–646.
- Clifford, P.; and Sudbury, A. 1973. A model for spatial conflict. *Biometrika*, 60(3): 581–588.
- Coates, A.; Han, L.; and Kleerekoper, A. 2018. A Unified Framework for Opinion Dynamics. In *International Conference on Autonomous Agents and MultiAgent Systems*, 1079–1086.
- Cooper, C.; Dyer, M.; Frieze, A.; and Rivera, N. 2018. Discordant voting processes on finite graphs. *SIAM Journal on Discrete Mathematics*, 32(4): 2398–2420.
- Cooper, C.; Elsasser, R.; Ono, H.; and Radzik, T. 2013. Coalescing random walks and voting on connected graphs. *SIAM Journal on Discrete Mathematics*, 27(4): 1748–1758.
- Cooper, C.; Radzik, T.; and Rivera, N. 2016. The coalescing-branching random walk on expanders and the dual epidemic process. In *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing*, 461–467.
- Corò, F.; Cruciani, E.; D’Angelo, G.; and Ponziani, S. 2022. Exploiting social influence to control elections based on positional scoring rules. *Information and Computation*, 289: 104940.
- Coste, S. 2021. The spectral gap of sparse random digraphs. *Ann. Inst. H. Poincaré Probab. Statist.*, 57(2): 644–684.
- Cruciani, E.; Mimun, H. A.; Quattropani, M.; and Rizzo, S. 2023. Phase transition of the k-majority dynamics in biased communication models. *Distributed Computing*, 36(2): 107–135.
- Cruciani, E.; Natale, E.; Nusser, A.; and Scornavacca, G. 2021. Phase transition of the 2-Choices dynamics on core-periphery networks. *Distributed Computing*, 34(3): 207–225.
- DeGroot, M. H. 1974. Reaching a consensus. *Journal of the American Statistical Association*, 69(345): 118–121.
- Del Vicario, M.; Bessi, A.; Zollo, F.; Petroni, F.; Scala, A.; Caldarelli, G.; Stanley, H. E.; and Quattrociocchi, W. 2016. The spreading of misinformation online. *Proceedings of the national academy of Sciences*, 113(3): 554–559.
- Donnelly, P.; and Welsh, D. 1983. Finite particle systems and infection models. *Mathematical Proceedings of the Cambridge Philosophical Society*, 94(1): 167–182.
- d’Amore, F.; Clementi, A.; and Natale, E. 2022. Phase transition of a nonlinear opinion dynamics with noisy interactions. *Swarm Intelligence*, 16(4): 261–304.
- d’Amore, F.; and Ziccardi, I. 2022. Phase transition of the 3-majority dynamics with uniform communication noise. In *International Colloquium on Structural Information and Communication Complexity*, 98–115. Springer.
- Faliszewski, P.; Gonen, R.; Koutecký, M.; and Talmon, N. 2022. Opinion diffusion and campaigning on society graphs. *Journal of Logic and Computation*, 32(6): 1162–1194.
- Friedkin, N. E.; and Johnsen, E. C. 1990. Social influence and opinions. *Journal of Mathematical Sociology*, 15(3-4): 193–206.

- Gärtner, B.; and Zehmakan, A. N. 2018. Majority Model on Random Regular Graphs. In Bender, M. A.; Farach-Colton, M.; and Mosteiro, M. A., eds., *LATIN 2018: Theoretical Informatics*, 572–583. Cham: Springer International Publishing. ISBN 978-3-319-77404-6.
- Hassin, Y.; and Peleg, D. 1999. Distributed probabilistic polling and applications to proportionate agreement. In *International Colloquium on Automata, Languages, and Programming*, 402–411. Springer.
- Hindersin, L.; and Traulsen, A. 2014. Counterintuitive properties of the fixation time in network-structured populations. *Journal of The Royal Society Interface*, 11(99): 20140606.
- Holley, R. A.; and Liggett, T. M. 1975. Ergodic theorems for weakly interacting infinite systems and the voter model. *The annals of probability*, 643–663.
- Krapivsky, P. L.; and Redner, S. 2003. Dynamics of majority rule in two-state interacting spin systems. *Physical Review Letters*, 90(23): 238701.
- Lesfari, H.; Giroire, F.; and Pérennes, S. 2022. Biased Majority Opinion Dynamics: Exploiting graph  $k$ -domination. In *International Joint Conference on Artificial Intelligence*.
- Liggett, T. M. 2012. *Interacting particle systems*, volume 276. Springer Science & Business Media.
- Lyons, R.; and Peres, Y. 2017. *Probability on trees and networks*, volume 42. Cambridge university press.
- Mobilia, M. 2003. Does a single zealot affect an infinite group of voters? *Physical review letters*, 91(2): 028701.
- Mobilia, M.; Petersen, A.; and Redner, S. 2007. On the role of zealotry in the voter model. *Journal of Statistical Mechanics: Theory and Experiment*, 2007(08): P08029.
- Montanari, A.; and Saberi, A. 2010. The spread of innovations in social networks. *Proceedings of the National Academy of Sciences*, 107(47): 20196–20201.
- Mountford, T.; Mourrat, J.-C.; Valesin, D.; and Yao, Q. 2016. Exponential extinction time of the contact process on finite graphs. *Stochastic Processes and their Applications*, 126(7): 1974–2013.
- Mourrat, J.-C.; and Valesin, D. 2016. Phase transition of the contact process on random regular graphs. *Electronic Journal of Probability*, 21: 1–17.
- Mukhopadhyay, A. 2024. Phase transitions in biased opinion dynamics with 2-choices rule. *Probability in the Engineering and Informational Sciences*, 38(2): 227–244.
- Mukhopadhyay, A.; Mazumdar, R. R.; and Roy, R. 2020. Voter and majority dynamics with biased and stubborn agents. *Journal of Statistical Physics*, 181: 1239–1265.
- Newman, M. 2018. *Networks*. Oxford university press.
- Out, C.; Tu, S.; Neumann, S.; and Zehmakan, A. N. 2024. The Impact of External Sources on the Friedkin–Johnsen Model. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management, CIKM '24*, 1815–1824. New York, NY, USA: Association for Computing Machinery. ISBN 9798400704369.
- Out, C.; and Zehmakan, A. N. 2024. Majority vote in social networks. *Information Sciences*, 657: 119970.
- Proskurnikov, A. V.; and Tempo, R. 2017. A tutorial on modeling and analysis of dynamic social networks. Part I. *Annual Reviews in Control*, 43: 65–79.
- Shirzadi, M.; and Zehmakan, A. N. 2025. Do Stubborn Users Always Cause More Polarization and Disagreement? A Mathematical Study. In *Proceedings of the Eighteenth ACM International Conference on Web Search and Data Mining, WSDM '25*, 309–317. New York, NY, USA: Association for Computing Machinery. ISBN 9798400713293.
- Sood, V.; Antal, T.; and Redner, S. 2008. Voter models on heterogeneous networks. *Physical Review E*, 77(4): 041121.
- Stewart, A. J.; Mosleh, M.; Diakonova, M.; Arechar, A. A.; Rand, D. G.; and Plotkin, J. B. 2019. Information gerrymandering and undemocratic decisions. *Nature*, 573(7772): 117–121.
- Van Der Hofstad, R. 2024. *Random graphs and complex networks*, volume 2. Cambridge university press.
- Wai, H.-T.; Scaglione, A.; and Leshem, A. 2016. Active Sensing of Social Networks. *IEEE Transactions on Signal and Information Processing over Networks*, 2(3): 406–419.
- Yildiz, E.; Ozdaglar, A.; Acemoglu, D.; Saberi, A.; and Scaglione, A. 2013a. Binary opinion dynamics with stubborn agents. *ACM Transactions on Economics and Computation (TEAC)*, 1(4): 1–30.
- Yildiz, E.; Ozdaglar, A.; Acemoglu, D.; Saberi, A.; and Scaglione, A. 2013b. Binary Opinion Dynamics with Stubborn Agents. *ACM Trans. Econ. Comput.*, 1(4).