

Opinion dynamics on random graphs Capannoli, F.

Citation

Capannoli, F. (2025, November 19). *Opinion dynamics on random graphs*. Retrieved from https://hdl.handle.net/1887/4283502

Version: Publisher's Version

Licence agreement concerning inclusion of doctoral

License: thesis in the Institutional Repository of the University of

Leiden

Downloaded from: https://hdl.handle.net/1887/4283502

Note: To cite this publication please use the final published version (if applicable).

CHAPTER 1 Introduction

§1.1 Opinion dynamics as interacting particles

The 1970s saw the birth of what would later be classified by Aldous [Ald13] as part of "the most broad-ranging currently active field of applied probability": interacting particle systems (IPSs). Such a persuasive description is justified by the interdisciplinary nature of the field, encompassing Mathematics, Physics, Computer Science, Biology, Social Sciences, and Epidemiology. In the mathematical community, foundational contributions to interacting particle systems were made by F. Spitzer [Spi69], who introduced models of infinite particle systems to explore how simple local interaction rules can generate intricate global large-scale stochastic behavior, and by R.L. Dobrushin [Dob71], who analyzed the existence and ergodicity of Markov processes with locally interacting components. After several years of intense developments, T. Liggett published a groundbreaking book in 1985 [Lig85] that became a cornerstone in the field. Liggett's work laid the foundation for a flourishing probabilistic theory, unifying diverse statistical mechanics phenomena under a single framework. His comprehensive approach synthesized rigorous results by using Markov process theory applied to phenomena such as magnetization, infection spread, and opinion dynamics. Additional references and modern introductions to the topic include [Dur88, Lig99, AF02, Swa22]. Models formally known as the stochastic Ising model, the contact process, and the voter model share a common feature: they can be represented as systems of locally interacting Markov chains labeled by a countable alphabet of types. The building block used to analyze such processes is the continuoustime random walk on graphs G = (V, E), having state space given by the vertex set V of the graph and infinitesimal generator given by

$$(\mathcal{L}_{rw})(f) = \sum_{x,y \in V} q(x,y) [f(x) - f(y)], \quad f: V \to \mathbb{R},$$
 (1.1)

where

$$q(x, y) = \frac{A(x, y)}{d_x}$$
 if $d_x > 0$, $q(x, y) = 0$ if $d_x = 0$, (1.2)

are the jump rates of the walk, A(x, y) denotes the adjacency matrix of the graph counting the number of edges $\{x, y\} \in E$, and d_x is the degree of vertex x, i.e., the number of edges $\{y, z\} \in E$ such that either y or z equals x. The rates q in (1.2) can be generalized to any infinitesimal generator representing a random walk. Informally, IPSs are multiple copies of the same process whose simple local interaction rules can lead to surprisingly rich and nontrivial dynamics. Initially, the central research question concerning interacting particle systems evolving over time was their long-term behavior; specifically, the existence and characterization of invariant measures (which are non-trivial in the infinite volume setting), and, whenever possible, the

determination of the domain of attraction of each such measure. The study of IPSs first focused on the lattice \mathbb{Z}^d , due to its regularity, translation invariance, and Euclidean geometric structure (through its embedding in \mathbb{R}^d). Within such a geometry, all aforementioned models exhibit phase transitions. In the stochastic Ising model and the contact process, phase transitions occur at a non-trivial critical value, depending on the dimension d, of the parameter that controls the intensity of magnetic interaction and infection transfer, respectively. Below this threshold, *fixation* takes place, that is, the global state of the particles remains the same although their local state may keep on changing, otherwise there is a positive probability that different states *coexist* indefinitely.

Interestingly, this dichotomy persists even in the opinion dynamics scenario, i.e. for the voter model, despite the lack of an intrinsic parameter that can be tuned. Here the phase transition is governed by the dimension d of the lattice: for $d \le 2$ fixation occurs, while for $d \ge 3$ coexistence emerges. Moreover, Holley and Liggett [HL75] proved the presence of a continuous family of invariant measures in the supercritical regime, with an index depending on the initialization of the opinions (state), while in the subcritical regime ($d \le 2$) every invariant measure reduces to a convex combination of Dirac measures at fixation states. These phenomena are closely related to the recurrence and transience properties of the simple random walk on \mathbb{Z}^d (see also [CG90, CGS95, CGS98]). In this thesis we focus on the study of the voter model on certain classes of finite graphs. In the following section we give a precise description of the model and some of its variants, and give a glimpse of some useful tools that we will exploit throughout the thesis.

§1.1.1 Voter model on fixed graphs

The voter model is a well-known example of opinion dynamics, introduced independently in the mathematical literature by Clifford and Sudbury [CS73], and by Holley and Liggett [HL75]. Like any IPS, it is defined by local update rules, which are specified by rates that depend on properties of the underlying graph, such as vertex degrees. In order to define the voter model in terms of its rates, we first require a suitable state space, and to this end we need a set of opinions, denoted by W. Let us assume W to be an alphabet, that is, a finite set with different elements. In this introduction, we restrict ourselves to the case |W| = 2 and, without loss of generality, let $W = \{0, 1\}$. These binary values, sometimes referred to as "blue" and "red," can be thought of as two competing political parties running in an election, hence the term "voter model". The question whether such a model can be adopted to give a statistical description of poll predictions has been discussed in the literature; see [FGSR+14] and references therein. Consider a locally finite graph G = (V, E), meaning that the graph itself can be infinite but each vertex has finite degree, and initialize the opinions according

to any distribution on $\Omega = \{0, 1\}^V$, the state space of the process. The evolution proceeds as follows: each vertex is equipped with an independent exponential clock of rate one. When its clock rings, it selects a neighbor uniformly at random and *adopts* that neighbor's opinion.

Formally, we define the voter model on G as the continuous-time Markov process $(\eta_t)_{t\geq 0}$ with state space Ω and infinitesimal generator \mathcal{L}_{vm} given by

$$(\mathcal{L}_{vm}f)(\eta) = \sum_{x \in V} \sum_{y \in V} q(x, y) \left[f(\eta^{x \to y}) - f(\eta) \right], \quad f : \{0, 1\}^V \to \mathbb{R},$$
 (1.3)

where $q(\cdot, \cdot)$ is given by (1.2) and

$$\eta^{x \to y}(z) = \begin{cases} \eta(y), & \text{if } z = x, \\ \eta(x), & \text{otherwise}. \end{cases}$$

For any $x \in V$ and $t \in \mathbb{R}^+$, $\eta_t(x)$ denotes the state of the vertex x at time t, taking values in the binary set $\{0, 1\}$, and interpreted as the opinion held by the individual x at time t, while $\eta_t = \{\eta_t(x) \mid x \in V\}$ is a configuration of opinions throughout the graph at time t. In other words, the process describes the evolution of opinion dynamics, starting from an initial configuration $\eta_0 = \{\eta_0(x) \mid x \in V\}$.

Notice that everything we have outlined can be generalized to any number of opinions $|W| \leq |V|$. Furthermore, the voter model considered here is the simplest model of opinion dynamics. Over the last few decades, many variants have been studied, and we now provide a brief overview. One of the most natural generalizations concerns the way opinions spread. In our version, the randomly selected vertex adopts the opinion of a random neighbor. This is also known as the pull voter model. Classical alternatives include the *push* voter model (or invasion process) [Dur10], in which the selected vertex imposes its opinion on a random neighbor, and the oblivious voter model, in which edges connecting vertices with differing opinions are sampled at random to drive the dynamics [CDFR18a]. Another important aspect is the nature of time evolution. In discrete time, the model can be either synchronous, where all vertices update their opinions simultaneously at each step, or asynchronous, where only one randomly selected vertex updates at a time. These differences can significantly affect the macroscopic behavior of the system. For a comprehensive introduction to these and related topics, we refer the reader to the PhD thesis of Nicolás Rivera [Riv18]. Further variants change the update rules more drastically and are of interest to the Mathematics, Computer Science, and Physics communities. The noisy voter model introduces spontaneous opinion changes (noise), modeling external influences and random fluctuations [GM95]. The q-voter model (or more generally, the threshold voter model) selects q neighbors (with replacement), and the vertex updates its opinion only if all q neighbors agree [DS93, CMnPS09]. A related variant is the k-majority

voter model, in which a vertex changes its opinion only if a majority of *k* sampled neighbors agree on the same opinion [dO92]. The final class of variants we discuss involves the introduction of a *bias* toward one or more opinions [HD19]. In this context, we mention the recent work on the majority voter model with a single bias [CMQR23], as well as our generalization analyzed in Chapter 4, where we study the behavior of the voting process under competing biases. Additional relevant contributions in applied contexts include [HP99, YPOS10, BCPS11].

We mainly study the classical voter model introduced above. Although we have implicitly assumed the graph G to be *undirected*, it is also possible to consider its directed counterpart. In the directed setting, each vertex $x \in V$ has both an in-degree d_x^- , representing the number of vertices with directed edges entering x, and an out-degree d_x^+ , representing the number of vertices with directed edges exiting x. The definition of the generator in (1.3) can be adapted by replacing d_x with d_x^+ in the expression for q(x, y). The definition of the random walk generator in (1.1) can be modified in the same way. For simplicity and introductory purposes, we will restrict our attention to the undirected version of the graph for the remainder of this section.

Consensus

Let us assume G to be *connected*, that is, given any pair of distinct vertices $x, y \in V$ there is a finite path (sequence of adjacent vertices) connecting them. Under this assumption, both the random walk and the voter model processes are irreducible Markov processes on G. Consider the voter model $(\eta_t)_{t\geq 0}$ on a finite $(|V| < \infty)$ connected graph G with state space $\Omega = \{0,1\}^V$. Then the process has two *absorbing states*, corresponding to the two monochromatic configurations, in which everyone share the same opinion, denoted by $\bar{1}$ and $\bar{0}$. In contrast to the infinite-volume setting of \mathbb{Z}^d , the voter model on a finite graph almost surely reaches an absorbing configuration in finite time. In other words, independent of the initial configuration the probability of having coexistence of opinions is zero. This setting naturally leads to the question of determining the time at which the system reaches the absorbing states, called *consensus time*, formally defined as

$$\tau_{\text{cons}} = \inf\{t \ge 0 : \eta_t \in \{\bar{1}, \bar{0}\}\}.$$
(1.4)

Note that this definition can be adapted to the case where G is disconnected. In this setting, τ_{cons} is defined as the first time at which all *connected components* reach a local consensus, which may differ across components. Moreover, the order of the consensus time will be determined by the time to local consensus within the largest connected component, possibly up to a constant factor depending on the graph in case this component is not unique. The analysis of (1.4) on finite connected graphs is, in full generality, far from complete. One reason is its strong dependence on the

underlying graph topology. A common approach to mitigate this complexity is to consider a graph sequence $\{G_n\}_{n\in\mathbb{N}}$, indexed by its size n=|V|, rather than analyzing a single fixed graph. In this context, an interesting interplay emerges between the voter dynamics and the volume n, leading to a central focus on the asymptotic behavior of the consensus time as $n\to\infty$. Understanding the asymptotic properties of $\tau_{\rm cons}$ on certain graph sequences will be one of the main objectives of this thesis.

In both recent and earlier literature, numerous works have aimed to establish sharp *bounds* on the expected consensus time $\mathbf{E}[\tau_{\text{cons}}]$, either in full generality (independent of the graph geometry) or for specific classes of graphs. Most approaches are based on the assumption that the graph is undirected, which ensures the reversibility of the associated random walk and enables the application of a variety of well-established techniques. Of particular importance is the result from [CEOR13a], where a spectral method is employed. More precisely, the authors show that under connectivity assumption,

$$\mathbf{E}[\tau_{\text{cons}}] = O\left(\frac{1}{(1 - \lambda_1)} \left(\log^4(n) + \frac{n}{\nu}\right)\right), \tag{1.5}$$

where λ_1 denotes the second largest eigenvalue of the transition matrix of the random walk on the graph, and ν is the ratio between the second moment of the degree sequence and the square of its first moment. This result gives crucial insight into a way to measure the speed of consensus. Namely, given a sharp estimate of the spectral gap $1 - \lambda_1$, we could tune ν by adjusting the *variability* of the degree sequence: increasing the variability decreases ν , thus speeding up consensus, and vice versa. In complete generality, Hassin and Peleg [HP99] proved that the dynamics reaches consensus on a state η_0 with probability proportional to the volume of vertices initially in state η_0 , in $O(n^3 \log n)$ rounds, regardless of the graph structure. This upper bound was recently improved to $O(n^3)$ by [KMS19], and was shown to be tight.

The first graph in which for the consensus time an exact asymptotic has been proved is the complete graph K_n , i.e., the graph in which each vertex is connected to all the other vertices. This was possible because the complete absence of geometry in the graph structure is a major simplifying feature. The voter model on K_n coincides with the so-called Moran model [Mor58], modeling the evolution of allele frequencies in a finite, constant-size population, capturing the effects of genetic drift, selection, and mutation. Using certain techniques that will be discussed in the next section, it is possible to reduce the problem to the analysis of simple random walks. This led to an explicit expression for the distribution of τ_{cons} . In particular,

$$\mathbf{E}[\tau_{\mathrm{cons}}] \sim n, \tag{1.6}$$

as $n \to \infty$. This example already provides insight into how long it takes to reach consensus on the complete graph K_n , despite its optimal connectivity properties. There are only two other cases for sequences of finite, undirected graphs in which the

exact asymptotic of $\mathbf{E}[\tau_{\text{cons}}]$ was understood: the d-dimensional tori by Cox [Cox89] and the random d-regular graph by Cooper and Frieze [CF04]. In Cox's work, the author showed that for $d \geq 3$ the expected consensus scales, up to an implicit prefactor, linearly in the size n of the graph, while for d = 1 and d = 2 the scaling is n^2 and $n \log(n)$, respectively. Similarly to the infinite lattice case, this dichotomy is closely related to the recurrence and transience properties of the random walk. Regarding the result by Cooper and Frieze, we first need to clarify the convergence claim, as two very dependent sources of randomness are active: one arising from the random geometry of the graph, the other from the stochastic process defined on it. We discuss this interplay in more detail in Section 1.2. For now, we note that the appropriate way to present their result is in the high-probability sense with respect to the randomness of the environment. That is, for any $\varepsilon > 0$,

$$\mathbb{P}\left(\left|\mathbf{E}[\tau_{\text{cons}}] - \frac{d-1}{d-2}n\right| > \varepsilon\right) \underset{n \to \infty}{\longrightarrow} 0, \tag{1.7}$$

where \mathbb{P} is the law of the random graph. Notice that both the d-dimensional tori and the random d-regular graph are symmetric, undirected and translation invariant geometries. In Chapter 2, we will prove what, to the best of our knowledge, is the first result for τ_{cons} for a wide class of *inhomogeneous directed* random graphs with fixed degree sequence. In the following section we discuss some of the most useful tools in the study of interacting particle systems, while describing how they can be exploited in the voter model context.

Graphical representation and duality

The expression of \mathcal{L}_{vm} in (1.3) indicates that the only possible transitions in an infinitesimal time step are of the form $\eta^{x\to y}$, meaning that only one opinion can change at a time. Consequently, the voter model belongs to the family of *spin systems*, a subclass of IPSs for which many general results have been proven; see [Lig85, Chapter 3] for further details. In particular, the voter model also falls into the class of attractive spin systems, that is, a site that takes a given value (0 or 1) is more likely to flip to the opposite value if it generally disagrees with its environment than if it generally agrees with it. Thus, the dynamics encourages agreement among neighboring sites. This property is essential because attractive spin systems are monotone, a feature that significantly facilitates their long-term analysis. A variety of tools have been developed for such IPSs, among which the most crucial and robust ones are graphical representation and duality. Indeed, rephrasing Liggett [Lig99, Part II], the main motivation for the introduction of voter models was their status as a class of spin systems for which the duality technique can be applied most completely and effectively. As we will argue below, many natural questions in this context reduce to problems involving systems of random walks, thereby revealing a deep connection

with one of the most active areas in probability theory over the past decades. What follows is inspired by the recent overview of interacting particle systems on random graphs in [CdH24, Chapter 3].

Duality. There are two ways to define a duality relation between two Markov processes $X = (X_t)_{t \ge 0}$ and $Y = (Y_t)_{t \ge 0}$, with state spaces X and Y, respectively: analytically or graphically. The graphical approach involves reversing time in the graphical representation of the process and using the evolution of the reversed process (which is often easier to analyze) to infer properties of the original process. The analytical approach requires identifying a bounded measurable function $H: X \times Y \to \mathbb{R}$ such that

$$E^{x}[H(X_{t}, y)] = E^{y}[H(x, Y_{t})],$$
 (1.8)

for any $(x, y) \in X \times \mathcal{Y}$ and $t \ge 0$. When this condition holds, we say that X and Y are dual to each other with respect to the duality function H. We first briefly describe the analytic duality and then the graphical duality, showing that they lead to the same result for the voter model.

Let $(A_t)_{t\geq 0}$ denote the dual Markov process with respect to the voter model $(\eta_t)_{t\geq 0}$, with state space $\Omega=W^{[n]}$, where |W|=2 are the number of opinions and $[n]=\{1,\ldots,n\},\,n\in\mathbb{N}$, is the vertex set V of the graph. Since we are considering a spin-flip system, the state space of the dual $(A_t)_{t\geq 0}$ is taken to be

$$\Xi = \{A : A \text{ is a finite subset of } [n]\},$$

which is finite for every $n \in \mathbb{N}$. Therefore, the dual process $(A_t)_{t\geq 0}$ is actually a Markov chain on Ξ , and can be interpreted as the time evolution of a collection of independent continuous-time random walks on [n] that coalesce whenever two of them occupy the same site, the so-called *coalescing random walk system*. It follows that $|A_t|$ can only decrease as t increases.

The duality function H that captures this behavior is

$$H(\eta, A) = \mathbf{1}_{\{\eta(x)=1 \ \forall x \in A\}}, \qquad \eta \in \Omega, \ A \in \Xi.$$

The main reason for this choice is that we are interested in consensus states, i.e., configurations η such that $\eta(x) = \eta(y)$ for all $x, y \in [n]$. Using this duality function together with the generator and the coefficients in (1.3), the transition rates of $(A_t)_{t\geq 0}$ for transitions $A \to B$, where $A, B \in \Xi$, are given by

$$Q(A,B) = \sum_{x \in A} \sum_{\substack{y \in [n]:\\ (A \setminus \{x\}) \cup \{y\} = B}} q(x,y).$$

This has the following interpretation: each $x \in A$ is removed from A at rate 1 and replaced by y with probability q(x, y). Moreover, if an attempt is made to place a point

at a site already occupied, then the two points coalesce. In this system of coalescing random walks, each walk independently waits an exponential time with rate 1, then moves according to q(x, y), and coalesces upon landing on an occupied site. The semigroup of each random walk is given by

$$P^{t}(x, y) = e^{-t} \sum_{n \in \mathbb{N}_{0}} \frac{t^{n}}{n!} q^{n}(x, y), \qquad x, y \in [n], \ t \ge 0.$$

All the previous statements can be generalised to a setting with |W| = k opinions with $k \in \{2, \ldots, n\}$ in the following way. The dual state space is taken to be $\Xi^{(k)} = \Xi^{k-1}$, the product of k-1 copies of Ξ . Consequently, the dual process is of the form $A_t^{(k)} = (A_1, \ldots, A_{k-1})_t$. We refer to López and Sanz [LS00] for an analytical description of duality for interacting particle systems. Here, the interpretation is slightly different from the case k=2. In the latter, $A \in \Xi$ represents the set of sites at which there are particles, while $A^{(k)} = (A_1, \ldots, A_{k-1}) \in \Xi^{(k)}$ represents the positions A_i of the particles that trace back to opinion i (this will become clear in the graphical representation), possibly with $A_i = \emptyset$ for some $i \in \{1, \ldots, k-1\}$. The dual function $H: \Omega \times \Xi^{(k)} \to \mathbb{R}$ now reads

$$H(\eta, A^{(k)}) = 1_{\{\eta(x)=i \ \forall \ x \in A_i \ \forall \ i \in \{1, \dots, k-1\}\}},$$

and the duality relation becomes

$$\mathbb{P}^{\eta} (\eta_t(x) = i \ \forall x \in A_i \ \forall i \in \{1, \dots, k-1\})$$
$$= \mathbb{P}^{A^{(k)}} (\eta(x_t) = i \ \forall x_t \in (A_i)_t \ \forall i \in \{1, \dots, k-1\})$$

for every initial configuration $\eta \in \Omega$ and every initial state $A^{(k)} \in \Xi^{(k)}$.

Let us consider the case in which k = n, and take $A^{(k)}$ such that $A_i = \{i\}$ for all sites, i.e., we place a continuous-time random walk at every site. Fix $\eta = \eta_0 \in \Omega$ to be the configuration in which each site has its personal opinion, say vertex $i \in [n]$ has opinion i, for which |[n]| = |W|. If one considers the configuration $\hat{\eta}_t$ defined as

$$\hat{\eta}_t(i) = \eta_0((A_i)_t), \qquad i \in [n], \ t > 0,$$

where $(A_n)_t = [n] \setminus \bigcup_{i \le n-1} (A_i)_t$,

then the duality relation says that $\hat{\eta}_t$ has the same distribution as the state η_t of the voter model at time t with initial configuration η_0 . Recall the definition of consensus time τ_{cons} with two opinions in (1.4), and define the *coalescence time* as

$$\tau_{\text{coal},n} = \inf\{t \ge 0 : \text{ all } n \text{ particles have coalesced into one}\}.$$

In particular, it follows that the consensus time has the same law as the coalescence time, specifically,

$$\mathbb{E}[\tau_{\text{cons}}] = \mathbb{E}[\tau_{\text{coal},n}].$$

Moreover, if we consider the same model with $2 \le k < n$ opinions and any given initial configuration η_0 then the distribution of $\tau_{\text{cons},k}$ will be *stochastically dominated* by the distribution of $\tau_{\text{coal},n}$ because $\{\tau_{\text{cons},k} > t\} \supseteq \{\tau_{\text{cons},n} > t\}$, so $\mathbb{P}^{\eta_0}(\tau_{\text{cons},k} \le t) \le \mathbb{P}(\tau_{\text{cons},n} \le t)$ for all $t \ge 0$. Here with $\tau_{\text{cons},k}$ we indicate the consensus time of the voter model with k opinions, with the convention that $\tau_{\text{cons},2} = \tau_{\text{cons}}$. Thus, for any $\eta_0 \in [k]^{[n]}$,

$$\tau_{\text{cons},k} \le \tau_{\text{coal},n}$$
, $2 \le k < n$, (1.9)

where \leq denotes the stochastic dominance of two random variables. In particular

$$\mathbb{E}^{\eta_0}[\tau_{\text{cons.k}}] \leq \mathbb{E}[\tau_{\text{coal.}n}], \qquad 2 \leq k < n.$$

Moreover, it can be proved (see e.g. Fernley and Ortgiese [FO23]) that, in the case where k = 2 and the initial distribution μ_u is given by the product measure of parameter $u \in (0, 1)$ Bernoulli random variables,

$$2u(1-u)\mathbb{E}[\tau_{\text{coal},n}] \leq \mathbb{E}^{\mu_u}[\tau_{\text{cons}}] \leq \mathbb{E}[\tau_{\text{coal},n}], \qquad u \in (0,1).$$

Graphical representation. We conclude this section by giving the duality principle in terms of the graphical representation. Start with the same setting as above: the voter model with state space $\Omega = W^{[n]}$, |W| = k, defined by its generator and its rates. Consider the graph $\{(j,t)\colon j\in [n], t\geq 0\}$ and independent rate-1 Poisson processes $(N_i(t))_{t\geq 0}, i\in [n]$. The dynamics is the following: if \bar{t} is an event of the clock N_i for some $i\in [n]$, then draw an arrow from (\bar{t},j) to (\bar{t},i) , where $j\in [n]$ is chosen with probability p(i,j). These transition probabilities coincide with the ones given above. In other words, an event represented by an arrow $j\to i$ means that at time \bar{t} the voter at site i decides to adopt the opinion of the voter at site j.

Given any initial configuration $\eta_0 \in \Omega$, we let the opinions flow upwards, starting at time t=0, and any time they encounter the base of an arrow they follow its direction, changing to the opinion that is at the tip of the arrow. In the case of two-opinions (0 and 1), this construction can be seen as a percolation process where a fluid is placed at t=0 in the 1-sites of η_0 and flows up the structure: the arrows are the pipes and the tips are the dams (see Durrett [Dur88]).

Let us next fix a time horizon $t_0 > 0$ and position a walk in (i, t_0) for all sites $i \in [n]$. We let these walks evolve independently as follows: they move downwards through the graph $\{(j, t_0 - t): j \in [n], t \in [0, t_0]\}$, and any time they encounter the tip of an arrow they follow it in the opposite direction. Furthermore, if one of them moves to a site already occupied by another walk, then the two walks coalesce into a single (independent) one. Alternatively, the process can be described as follows: each of the walks waits an exponential time of parameter 1 and, given the current position $x \in [n]$, moves to $y \in [n]$ with probability q(x, y). Likewise, for the coalescing condition we have that any time when at least two walks meet at the same site they

coalesce into a single one. Denote by $A_t^{(n),t_0} = (A_1^{t_0}, \dots, A_n^{t_0})_t$ the resulting system of n coalescing random walks (CRWs) evolving as above, where, for each $i \in [n]$, $(A_i)_t^{t_0}$ is the position of the walk starting in (i,t_0) at time t, in particular, $A_0^{(n),t_0} = [n]$.

Given this construction, it follows that the opinion held by vertex i at time t_0 can be derived by *tracing back in time* the path of the walks up to time t = 0. Thus,

$$\eta_{t_0}(i) = \eta_0((A_i)_{t_0}^{t_0}) \qquad \forall i \in [n], \, \forall t_0 > 0.$$

With the latter we derived the same result obtained by using the duality relation. All the other results regarding the equivalence in distribution between the coalescing time and the consensus time with k = n opinions, i.e., $\eta_0 = [n]$, follow directly.

In modern probability theory, analytic duality stands out as one of the most powerful and versatile tools, playing a pivotal role in cutting-edge research. Numerous studies have significantly advanced this theory. Among them, notable contributions have been provided by G. Carinci, F. den Hollander, C. Franceschini, C. Giardiná, S. Nandan, E. Pulvirenti, F. Redig, F. Sau, T. Sasamoto, and J. Swart [RS18, dHN22, CGRS16, BdHM⁺24, FG19].

§1.2 Random graphs models

Initially, the only finite graph sequences that were analytically tractable and extendable were the complete graph K_n and the d-dimensional tori. The latter preserves the Euclidean geometry, when embedded in \mathbb{R}^d , of the lattice \mathbb{Z}^d within a finite-volume context, maintaining translation invariance. In contrast, the complete graph K_n represents the opposite extreme, where geometric structure is entirely absent. Both applied and theoretical sciences have required new finite models capable of describing the complexity involved in the evolution of opinion dynamics and related models on more realistic approximations of real-world networks. This need was met with the emergence of $random\ graph$ theory.

Formally, a random graph is a graph-valued random variable, where the randomness may arise from different sources: either by fixing the vertices and randomizing the edges, randomizing both the edges and the vertices, or fixing both the vertices and their degrees, with the randomness lying in how the edges are formed under these constraints. Similarly to what was described in Section 1.1, it will also be necessary in this framework to consider (random) graph sequences $\{G_n\}_n$ in order to study asymptotic properties A_n , indexed by the size n, as n grows to infinity. In particular, we say that a property A_n holds with high probability if $\mathbb{P}(A_n) \to 1$ as $n \to \infty$, where $\mathbb{P} = \mathbb{P}_n$ is the law of G_n . In this view, over the past couple of decades, several topologies have been introduced that allow for a detailed study of the asymptotic behavior of graph sequences $\{G_n\}_n$. The most commonly used are local weak convergence [BS01, AS04]

and *graphons* [LS06]. The latter is used to study the asymptotic properties of *dense* random graphs, that is, graph sequences in which the number of edges is of the same order as the number of vertices, and consists of a symmetric, measurable function from the unit interval [0, 1] to itself equipped with a specific metric. In contrast, local weak convergence is suitable for *sparse* random graphs, where the number of edges is of the same order as the volume, and provides information about the structure of the graph based on a local exploration process driven by its composition. In this section, we will describe some classes of *sparse* random graph models, focusing on the so-called *configuration model*. We refer to [Bol01, FK15, Hof17, Hof24] for comprehensive references on random graph theory.

Independent edges random graphs

The earliest random graph models are referred to as Erdös-Rényi random graphs. There are two distinct models commonly known by this name. The first is the Gilbert-Erdös-Rényi random graph [Gil59], denoted by G(n, p). This random graph model can be sampled as follows: initially take an empty graph (no edges) with $n \in \mathbb{N}$ vertices, and for each unordered pair of vertices $x, y \in [n]$, add an edge with probability $p \in [0, 1]$, possibly depending on the size n of the graph. With probability 1 - p, no edge is added. This procedure is equivalent to performing bond percolation on the complete graph K_n . The second model is the Erdös-Rényi random graph [ER59], denoted by G(n, M), where a graph is chosen uniformly at random from the set of all graphs with n vertices and M edges. The key difference between G(n, p) and G(n, M) is that in G(n, p) edges are included independently with fixed probability p, whereas G(n, M) has a fixed total number of edges. In the latter, edges are not independent; for example, if one edge is present, only M-1 edges remain to be assigned among the remaining vertex pairs. Despite this difference, the two models are closely related. In fact, conditioning G(n, p) on having exactly M edges yields the uniform distribution over graphs with M edges, which is precisely G(n, M). Conversely, G(n, M) can be viewed as G(n, p) conditioned on a fixed edge count M. Thus, many properties of one model can be translated to the other via conditioning or averaging arguments. Indeed, G(n, p) and G(n, M) are asymptotically equivalent for a wide range of graph properties when M and p are appropriately matched. If $M = \binom{n}{2}p + O(n\sqrt{p(1-p)})$, i.e., if M is close to its mean under G(n, p), then any property P that holds with high probability in G(n, M) also holds with essentially the same probability in G(n, p) as $n \to \infty$ (see e.g. [FK15]).

Due to this equivalence, it is often sufficient to study just one of the two models. Let us focus on G(n, p). More precisely, consider the *sparse* regime of the model, where the (expected) number of edges is of the same order as the number of vertices.

In other words, we assume that

$$p=\frac{\lambda}{n}$$

where $\lambda > 0$ does not depend on n. From a probabilistic perspective, the model basically reduces to a study of binomial random variables and, despite its simplicity, interesting phase transitions phenomena emerge. More precisely, it has been proved that $\lambda = 1$ is the critical value for a drastic change in the order of the largest connected component (LCC). Indeed, if $\lambda < 1$, then with high probability the size of the LCC scales as $\log(n)$; if $\lambda = 1$ then the LCC scales as $n^{2/3}$; while if $\lambda > 1$, then the LCC is unique and is a *giant* component, i.e. its size scales as n, while all other connected components have size $O(\log(n))$ with high probability. Moreover, if $p > C\frac{\log(n)}{n}$ for some C > 1, then the resulting graph will be connected with high probability. Many more refined results regarding the fine structure of this random graph ensemble is known; we refer to [Hof17, Chapters 3 & 4] for details.

Although Erdös-Rényi random graphs marked the beginning of a conceptual revolution, they exhibited a clear limitation in terms of homogeneity: they were not realistic models of real-world complex networks. Building on this observation, new models with diverse features were introduced. One of the key models in this direction is the *inhomogeneous* Erdös-Rényi model, also known as the generalized random graph, introduced in [BJR07]. Here, the construction is similar to G(n, p), but instead of having the same p for each pair of vertices $x, y \in [n]$, the connection probabilities may differ depending on the selected vertices. More precisely, we consider a weight sequence $\{w_x\}_{x\in[n]}$, possibly depending on n, which can either be fixed or be sampled according to independent and identically distributed random variables, and we define

$$p_{x,y} = \frac{w_x \, w_y}{\sum_z w_z + w_x \, w_y}$$

to be the independent connection probability between any pair of vertices x, y. Note that, by taking $w_x = \frac{n\lambda}{n-\lambda}$, we obtain $p_{x,y} = p = \frac{\lambda}{n}$ for any $x, y \in [n]$, thus recovering the classical G(n,p) case. Consider the case in which $\{w_x\}_{x\in[n]}$ is a deterministic sequence, and let D_n denote the degree of a uniformly selected vertex. Under certain assumptions, particularly regarding the convergence in distribution of D_n to some limiting random variable D, along with convergence of its first and second moments, many structural properties of this ensemble can be established. In particular, the explicit limiting distribution of the degree sequence is known, and is given by a mixed Poisson distribution with the mixing distribution having the law of D. Several other inhomogeneous edge-independent models have been proposed based on a weight sequence and different specifications for the connection probabilities $p_{x,y}$. Among the most studied are the Chung–Lu [CL02b, CL02a] and Norros–Reittu [NR06] random graphs. Under suitable assumptions on the weight sequence, it can be shown that

all these models are asymptotically equivalent (see [Hof17, Chapter 6]). Recent developments in the study of stochastic processes in these geometries include [ST20, BP25, BLPS18, AGSS24]. Furthermore, there has been an increasing interest in the study of spectral properties of inhomogeneous random graphs [BL10, CHdHS21, CHMS25]. We refer to [HHM23b, HHM23a] for a detailed analysis of the local weak convergence of inhomogeneous random graphs.

Another direction in the generalization of Erdös-Rényi random graphs involves embedding the vertices in a metric space and adjusting the connection probabilities $p_{x,y}$ by introducing a parametrized function that modulates the likelihood of connecting two vertices based on their distance. Typically, this results in taking the inverse of the distance between x and y, raised to some power $\gamma > 0$. These models are known as *geometric random graphs* (see [Pen03, DDC23] for more details). Finally, there is another random graph ensemble that does not fall into the class of models with independent edges but is nevertheless worth mentioning: the *preferential attachment* random graph. This model, introduced by Barabási and Albert [BA99], describes the growth of networks in which new vertices prefer to connect to existing vertices with higher degrees. This mechanism leads to the emergence of *scale-free networks*, characterized by a power-law degree distribution, where a few vertices (hubs) accumulate many connections, while the majority of vertices have only a few.

§1.2.1 Configuration models

In complex network theory, one of the key elements in determining the flow of information diffusion along a network is the inhomogeneity of the *degree sequence*. There was a need to study opinion dynamics on empirical networks where the degree sequence of the vertices was known, but the way the connections were distributed along the network was too complicated to analyze case-by-case. The *configuration model* was the right tool to address this problem. It was first introduced by Bollobás in 1980 [Bol80] to study the enumeration of regular graphs, and was later extended to general degree sequences by Molloy and Reed in 1995 [MR95], who analyzed the emergence of a giant component and established critical thresholds for connectivity.

The configuration model is sampled as follows. For any $n \in \mathbb{N}$, let $[n] := \{1, \ldots, n\}$ be a set of n labeled vertices. Define $\mathbf{d}_n = (d_1, \ldots, d_n)$ to be a deterministic degree sequence satisfying the following constraint

$$M_n := \sum_{x \in [n]} d_x = 2 \,\ell_n \,, \tag{1.10}$$

for some $\ell_n \in \mathbb{N}$. The randomness of the model comes from the mechanism in which the edges are formed, and is the result of the following uniform pairing procedure involving *stubs*, i.e., half-edges. Assign to each vertex $x \in [n]$, d_x labeled stubs. At

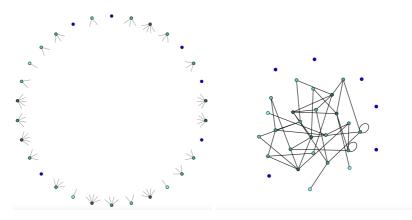


Figure 1.1: On the left the initial empty graph with all vertices having stubs according to their degree. On the right a possible realization of a configuration model with the prescribed degree sequence. The graph is created using software available at https://www.networkpages.nl.

each step, select a stub e that was not matched in a previous step (the order of the selection of e is irrelevant), and a uniform at random stub f among the unmatched ones. Then match them and add the edge $ef \in E$ between the vertex incident to e and the vertex incident to f to the edge set f. Continue until there are no more unmatched stubs. See Figure 1.1 for a possible realization following such a procedure. Note that the matching is possible by the assumption (1.10) that the sum of the stubs f0 is even. This random procedure gives rise to a so-called f1 configuration, and it uniquely determines the corresponding random graph f2 configuration, with a given degree sequence f3 it is sampled according to the procedure above. Notice that the pairing procedure is well-defined and can be extended to degree sequences sampled according to i.i.d. random variables.

There are two main differences between the generation of $CM(\mathbf{d}_n)$ and any of the graphs mentioned in the previous subsection: in the configuration model, the connection probabilities are *not independent*, and the resulting graph is *not necessarily simple*, meaning that there may be self-loops (edges from a vertex to itself) and multiple edges (more than one edge connecting the same pair of vertices). Under certain conditions on the degree sequence \mathbf{d}_n , the explicit distribution of the number of self-loops and multiple edges is known. Interestingly, under the same assumptions, the general properties of $CM(\mathbf{d}_n)$ and those of the same ensemble with self-loops and multiple edges erased are asymptotically equivalent. Another important property is that $CM(\mathbf{d}_n)$ is not uniformly distributed among all multigraphs (i.e., not necessarily simple graphs) with degree sequence \mathbf{d}_n . This is due to the fact that the half-edges attached to the same vertex are indistinguishable. Nevertheless, if we condition the graph on being simple, then the resulting law of $CM(\mathbf{d}_n)$ becomes uniform over all

simple graphs with degree sequence \mathbf{d}_n . See [Hof17, Chapter 7] for further details.

Directed model

The directed configuration model, whose structural properties were first studied by Cooper and Frieze [CF04], is a natural generalization to the directed setting of the classical version of the model defined in the previous section. The sampling procedure for this random graph ensemble is similar to the undirected counterpart and can be described as follows. In the same setting as above, let $\mathbf{d}^+ = \mathbf{d}_n^+ = (d_x^+)_{x \in [n]} \in \mathbb{N}_0^n$ and $\mathbf{d}^- = \mathbf{d}_n^- = (d_x^-)_{x \in [n]} \in \mathbb{N}_0^n$ be two finite deterministic sequences such that

$$m = m_n := \sum_{x \in [n]} d_x^+ = \sum_{x \in [n]} d_x^-. \tag{1.11}$$

For any vertex $x \in [n]$, let d_x^+ (resp. d_x^-) be its out-degree (resp. in degree), that is, the number of vertices that are connected to x via a directed edge that is exiting (resp. entering) x. We perform a pairing procedure of the same type as for the undirected model, but this time we need to match in-stubs, called *heads*, with out-stabs, *tails*. At each step, select a tail e that was not matched in a previous step, and a uniform at random head f among the unmatched ones, match them, and add the directed edge ef between the vertex incident to e and the vertex incident to f to the edge set f. Continue until there are no more unmatched heads and tails. As before, note that the constraint in f (1.11) ensures that such a uniform matching ends without any stub left unmatched. We say that a graph f is sampled from the directed configuration model DCM = DCM(\mathbf{d}_n^+ , \mathbf{d}_n^-) with a given bi-degree sequence \mathbf{d}_n^+ , \mathbf{d}_n^- if it is sampled according to the procedure above.

There are many differences between the configuration model (CM) and its directed counterpart (DCM), and more generally between directed and undirected (random) graphs. Notice first that the law of the random walk on any random graph is a random measure. In the case of the CM, the stationary distribution π for the random walk is, despite the graph being random, deterministic and explicit, given by

$$\pi(x) = \frac{d_x}{\sum_z d_z},$$

and the random walk is reversible with respect to π . In contrast, in the DCM, the stationary distribution π is a non-trivial and, in most cases, *non-reversible* random measure. This immediately implies that many of the techniques used to study these processes in the undirected case cannot be applied in the directed setting. Nevertheless, in the last decade, new tools have emerged that allow for a detailed study of this topic. Within this domain, the works of C. Bordenave, X. S. Cai, P. Caputo, G. Perarnau, M. Quattropani, and J. Salez stand out as particularly significant [CP21, CCPQ23,

BCS19, BCS18, CP20]. Under various assumptions on the degree sequence $(\mathbf{d}_n^+, \mathbf{d}_n^-)$ of the DCM, they proved several results concerning the *uniqueness* of π , its *extremal values* π_{max} , π_{min} , its *bulk* structure, and the *mixing time* t_{mix} , with high probability. In particular, in [BCS18] the authors show that the mixing time is logarithmic in n and that the total variation distance drops abruptly to zero at a sharp time, an instance of the so-called *cutoff at the entropic time* (see also [BCS19, CCPQ23]).

In this thesis, we will primarily study the voter model evolving on the DCM, focusing on the *sparse* version of the model, that is, working under assumptions on the degree sequence such that $m_n = \Theta(n)$. Under these assumptions, it can be shown that the local weak limit is an explicit Galton–Watson tree. One consequence of this result is that a local exploration process, such as the breadth-first search from a randomly selected node up to a certain depth, will with high probability not encounter any cycles. This property is crucial since, due to the duality relation between the voter model and a system of coalescing random walks (see Section 1.1) and the groundbreaking results of Oliveira [Oli12, Oli13], it is possible to reduce the study of the consensus time τ_{cons} to the analysis of the meeting time of two random walks.

Voter model on random graphs

In recent years, there has been increasing interest in the study of the voter model, mainly through the analysis of consensus time, on various underlying random geometries. In [FO23], the authors provide explicit estimates, up to polylogarithmic corrections, of the expected consensus time for a wide range of independent-edges, subcritical (low edgedensity regime) random graphs, while [Fer24] explores supercritical random graphs. A detailed analysis of sparse d-regular random graphs is presented in [ABH⁺24], where the authors build on the works [CF05, CFR10], investigating the evolution of discordant edges over different time scales. There are also numerous results concerning observables of random walks on such geometries, which implicitly yield results for the voter model via duality. In this direction, we highlight the work by Oliveira [Oli12, Oli13], which extends the so-called *mean-field* behavior known for the complete graph K_n to general classes of (random) graphs, as well as the work by Chen, Choi, and Cox [CCC16] and Chen [Che18], which offers key insights into martingale properties of voter density processes, leading to convergence to the Wright-Fisher diffusion. Furthermore, Hermon et al. [HLYZ22] analyze the decay of the particle density in the coalescing random walks model during the so-called "Big Bang" regime, that is, an early phase of the process when many particles are present and interactions are frequent. A spectral approach was employed in [BK24] to compute the expected meeting time from stationarity for two random walks on certain classes of independentedges random graphs. We also mention that dynamic random graph models, in which the voter evolution affects the evolution of the underlying graph, have recently

received attention. Various rewiring rules coupled with voter dynamics were proposed in [DGL+12], with predictions based on numerical simulations. Recent works such as [BS17, AdHR25, BBdHM24, BK25] have rigorously studied these models in the setting of dense random graphs, while [ABH+25] considers the sparse regime. Lastly, we cite the updated version of Durrett's book in progress [Dur25], the survey article by Capannoli and den Hollander [CdH24], and the references mentioned therein, for an overview of recent achievements on this topic. In the following section, we provide a brief description of how our contributions relate to the current literature and outline possible future directions for this exciting field of research.

§1.3 Overview

In this section we present a brief overview of the content of the four chapters in this thesis.

I. In Chapter 2, we consider Markovian dynamics on a typical realization of the directed configuration model. In this random geometry, we study the consensus time of the voter model on a typical realization of the graph, as well as the meeting time of two random walks starting from stationarity. It is known that these two quantities are related when the underlying sequence of graphs satisfies certain mean-field conditions. Previous results in this direction were restricted to settings with prominent symmetry properties in the environment, such as the complete graph, discrete d-dimensional tori ($d \ge 2$), and random regular graphs. To the best of our knowledge, our work is the first to address this problem in a general random directed geometry.

We provide a complete characterization of the distribution of the meeting and consensus times on a typical random graph as a function of a single quantity ϑ . More precisely, we show that for a typical large graph drawn from the DCM ensemble, the distribution of the meeting time is well approximated by an exponential random variable. Furthermore, we provide the precise first-order approximation of its expectation, showing that it is linear in the size of the graph. The explicit constant ϑ depends on simple statistics of the degree sequence. As a consequence, we can analyze the effect of the degree sequence on the expected meeting time and, through explicit examples, examine how its regularity or variability plays a role in the diffusion of opinions. Finally, we perform a quantitative analysis of ϑ , addressing how degree variability and correlations between in- and out-degrees influence the consensus time, and identifying conditions that minimize or maximize it under structural constraints.

II. After gaining such an explicit result in Chapter 2, we proceed to extend it to a deeper level of analysis for the voter model dynamics. In Chapter 3, we

investigate the evolution of the voter model on sparse directed random graphs generated via the DCM, focusing on the density of discordant edges, i.e., edges connecting vertices with opposite opinions. This type of analysis was performed earlier in the undirected regular random graph setting.

Exploiting duality and developing new coupling techniques, we derive precise asymptotics for the expected density of discordant edges across different time scales, capturing not only the consensus time but also the transient behavior of the dynamics. More precisely, we prove the existence of an initial time scale in which the process behaves as if it were evolving on a Galton-Watson tree, for which the initial density of discordances quickly drops to a given state. For time scales strictly in between this initial drop and the consensus time, a quasistationary behavior emerges and the density of discordant edges stabilizes around an explicit limiting value for a long time. Finally, on the consensus time scale, the process exhibits a sharp descent from the previous plateau, approaching zero at a rate that coincides with the functional ϑ found in the previous chapter for the expected consensus time.

III. In Chapter 4, we analyze a nonlinear opinion dynamics model on sparse random directed graphs, designed to capture the interplay between an external disruptive bias favoring the adoption of a novel (blue) opinion and individual stubbornness reinforcing adherence to the initial (red) opinion. Our model introduces two competing parameters: a bias probability *p*, modeling external influences, and a stubbornness parameter *s*, representing intrinsic resistance to opinion change. We first analyze the model on random *d*-out-regular directed graphs, initially configured with all agents supporting the red opinion.

Utilizing a mean-field approximation, in which network topology plays no explicit role and the underlying graph is independently resampled at each step, we identify a phase transition phenomenon occurring with high probability. For fixed out-degree $d \geq 2$ and stubbornness s > 1, we establish the existence of a critical threshold $p_c(d,s)$ of the disruptive bias that separates two distinct dynamical behaviors. Above this threshold $(p > p_c(d,s))$, the dynamics quickly reaches consensus on the blue opinion. Below the threshold $(p < p_c(d,s))$, the system rapidly settles into a metastable state characterized by a stable nonconsensus proportion $q_{\star}(p,d,s) < 1$ of agents adopting the blue opinion. This metastable regime persists on a time scale that is exponential in the network size. We prove these results for both synchronous dynamics, where all agents update simultaneously, and asynchronous dynamics, where only one randomly selected agent updates at each time step. Furthermore, we perform numerical simulations on sparse homogeneous and inhomogeneous random graphs, to confirm that the qualitative and quantitative predictions from our mean-field setting can

be extended to more generic frameworks.

IV. In Chapter 5, we investigate several questions related to the voter model on both directed and undirected random graph models, particularly focusing on heavy-tailed degree distributions. While various predictions for the consensus time order exist in the physics literature for undirected random graphs with uncorrelated degrees, we conduct a systematic numerical study and propose multiple conjectures based on mathematical reasoning.

Our analysis in Chapter 5 explores how the network topology, especially heterogeneity in vertices degrees, affects the consensus time. The main contribution is an extension of the analysis of consensus times to directed random graphs with arbitrary given degree sequences. This work generalizes the analytical results derived in Chapter 2, initially established for bounded deterministic degrees, to a broad range of directed random graph ensembles. Specifically, we test conjectures numerically for degrees following a Pareto distribution with exponent $\alpha > 0$, and we derive precise first-order asymptotics for the expected consensus time across all values of $\alpha > 0$. Additionally, we evaluate the validity of mean-field approximations within these graph ensembles. Numerical simulations demonstrate that mean-field approximations accurately describe the distribution of $\tau_{\rm cons}$ provided the underlying directed ensemble has finite mean degrees ($\alpha > 1$). Conversely, in scenarios with infinite mean degrees ($\alpha \le 1$), significant deviations occur between the empirical distribution of consensus times and the theoretical predictions. Lastly, we analyze a similar crossover regarding the convergence of a specific weighted density process to a Wright-Fisher diffusion. When the underlying graph has finite mean in- and out-degrees, we observe convergence for both directed and undirected configuration models. In the directed model, we explicitly identify the diffusion constant as a function of the in- and out-degree distributions.

Future directions

As outlined throughout this thesis, over the past few years we have developed a significantly deeper understanding of how consensus is reached in various opinion dynamics models evolving on sparse directed and undirected random geometries. Our contributions span both theoretical advancements and applied insights into the behavior of such processes.

The field remains rich with opportunities for further exploration. First and fore-most, the heuristic phenomena presented in Chapter 5 await rigorous mathematical validation. In addition, numerous extensions of the voter model on random graphs introduce new layers of complexity and open up various new research directions. These include, for instance, the study of *dynamic* random graphs and *feedback-driven dynamic* random graphs are feedback-driven dynamic.

ics (for recent works in this direction see, e.g. [ST20, PSS20, dSOV24, AvdHdHN25, ABH+25, SV25] and [AdHR25, BBdHM24], respectively) both in directed and undirected sparse and dense settings. Beyond that, variants of the voter model itself, such as the version incorporating competing biases discussed in Chapter 4, present further challenges and opportunities for advancing our understanding of opinion dynamics in complex networks.