On the Role of Memory in Robust Opinion Dynamics

Luca Becchetti[1] , **Andrea Clementi**[2] , **Amos Korman**[3] , **Francesco Pasquale**[2] , **Luca Trevisan**[4] and **Robin Vacus**[5]

[1]*Sapienza* University of Rome, Rome, Italy
[2]*Tor Vergata* University of Rome, Rome, Italy
[3] CNRS, located at FILOFOCS, Tel-Aviv, Israel
[4]*Bocconi* University, Milan, Italy
[5]CNRS, IRIF, Paris, France
becchetti@diag.uniroma1.it, clementi/pasquale@mat.uniroma2.it, amos.korman@irif.fr,
l.trevisan@unibocconi.it, rvacus@irif.fr

Abstract

We investigate opinion dynamics in a fully-connected system, consisting of n agents, where one of the opinions, called *correct*, represents a piece of information to disseminate. One *source* agent initially holds the correct opinion and remains with this opinion throughout the execution. The goal of the remaining agents is to quickly agree on this correct opinion. At each round, one agent chosen uniformly at random is *activated*: unless it is the source, the agent pulls the opinions of ℓ random agents and then updates its opinion according to some rule. We consider a restricted setting, in which agents have no memory and they only revise their opinions on the basis of those of the agents they currently sample. This setting encompasses very popular opinion dynamics, such as the *voter model* and *best-of-k majority* rules.

Qualitatively speaking, we show that lack of memory prevents efficient convergence. Specifically, we prove that any dynamics requires $\Omega(n^2)$ expected time, even under a strong version of the model in which activated agents have complete access to the current configuration of the entire system, i.e., the case $\ell = n$. Conversely, we prove that the simple voter model (in which $\ell = 1$) correctly solves the problem, while almost matching the aforementioned lower bound.

These results suggest that, in contrast to symmetric consensus problems (that do not involve a notion of correct opinion), fast convergence on the correct opinion using stochastic opinion dynamics may require the use of memory.

1 Introduction

Identifying the specific algorithm employed by a biological system is extremely challenging. This quest combines empirical evidence, informed guesses, computer simulations, analyses, predictions, and verifications. One of the main difficulties stems from the huge variety of possible algorithms, which is particularly true when multi-agent systems are concerned [Sumpter, 2010; Feinerman and Korman, 2017b]. To reduce the space of algorithms, the scientific community often restricts attention to simple algorithms [Couzin *et al.*, 2005; Gelblum *et al.*, 2015; Fonio *et al.*, 2016]. However, even though this restriction reduces the space of algorithms significantly, the number of simple algorithms still remains extremely large.

Another direction to reduce the parameter space is to disqualify certain algorithms because they are unable to efficiently handle the challenges induced by the corresponding scenario [Boczkowski *et al.*, 2018; Feinerman and Korman, 2017a; Bialek *et al.*, 2012]. Analyzing the limits of computation has been a main focus in theoretical computer science. Hence, it appears promising to employ lower-bound techniques from this discipline to biologically inspired contexts, in order to identify which algorithms are less likely to be used, or alternatively, which parameters of the setting are essential for efficient computation [Guinard and Korman, 2021; Boczkowski *et al.*, 2018]. This lower-bound approach may help identify and characterize phenomena that might be harder to uncover using more traditional approaches, e.g., using simulation-based approaches or differential equations techniques. The downside of this approach is that it is limited to analytically tractable settings, which may be too "clean" to perfectly capture a realistic setting.

Taking a step in the latter direction, we focus on a basic problem of information dissemination, in which few individuals have pertinent information about the environment, and other agents wish to learn this information despite having constrained and random communication [Aspnes and Ruppert, 2009; Boczkowski *et al.*, 2019; Bastide *et al.*, 2021; Korman and Vacus, 2022]. Such information may include, e.g., knowledge about a preferred migration route [Franks *et al.*, 2002; Lindauer, 1957], the location of a food source [Couzin *et al.*, 2011], or the need to recruit agents for a particular task [Razin *et al.*, 2013]. In some species, specific signals are used to broadcast such information, a remarkable example being the waggle-dance of honeybees [Franks *et al.*, 2002; Seeley, 2003]. In many other biological systems, however, it may be difficult for individuals to distinguish those who have pertinent information from others in the group [Couzin *et al.*, 2005; Razin *et al.*, 2013]. Moreover, in multiple biological contexts, animals cannot rely on distinct signals and must obtain information by merely observing the behavior of other animals (e.g., their position in space, speed, etc.). This weak form of communication, often referred to as *passive communication* [Wilkinson, 1992], does not even require animals to deliberately send communication signals [Cvikel *et al.*, 2015; Giraldeau and Caraco, 2018]. A key theoretical question is identifying minimal computational resources that are necessary for information to be disseminated efficiently using passive communication.

Here, following the work in [Korman and Vacus, 2022], we consider an idealized model, that is inspired by the following scenario.

Animals by the pond. Imagine n animals that gather around a pond to drink water from it. Assume that one side of the pond, either the northern or the southern side, is preferable (e.g., because the risk of having predators there is reduced). However, the preferable side is known to a few animals only. These informed animals will remain on the preferable side of the pond. The remaining group members would like to learn which side of the pond is preferable, but they are unable to identify the knowledgeable animals. Instead, they can scan the pond and estimate the number of animals on each side of it, and then, according to some rule, move from side to side. Roughly speaking, the main result in [Korman and Vacus, 2022] is that there exists a rule that allows all animals to converge on the preferable side quickly, despite initially being spread arbitrarily in the pond. The suggested rule essentially says that each agent compares its current sample with the sample obtained in the previous round; If it sees that more animals are on one particular side now than they were in the previous sample, then it moves to that side.

Within the framework described above, we ask whether knowing anything about the previous samples is really necessary, or whether fast convergence can occur by considering the current sample alone. Roughly speaking, we show that indeed it is not possible to converge fast on the correct opinion without remembering information from previous samples. Next, we describe the model and results in a more formal manner.

Problem definition. We consider n agents, each of which holds an *opinion* in $\{0, 1, \ldots, k\}$, for some fixed integer k. One of these opinions is called *correct*. One *source* agent[1] knows which opinion is correct, and hence holds this opinion throughout the execution. The goal of non-source agents is to converge on the correct opinion as fast as possible, from any initial configuration. Specifically, the process proceeds in discrete *rounds*. In each round, one agent is sampled uniformly at random (u.a.r) to be *activated*. The activated agent is then given access to the opinions of ℓ agents, sampled u.a.r (with replacement[2]) from the multiset of all the opinions in the population (including the source agent, and the sampling agent itself), for some prescribed integer ℓ called *sample size*. If it is not a source, the agent then revises its current opinion using a decision rule, which defines the *dynamics*, and which is used by all non-source agents. We restrict attention to dynamics that are not allowed to switch to opinions that are not contained in the samples they see. A dynamics is called *memoryless* if the corresponding decision rule only depends on the opinions contained in the current sample and on the opinion of the agent taking the decision. Note that the classical *voter model* and *majority* dynamics are memoryless.

1.1 Our Results

In Section 3, we prove that every memoryless dynamics must have expected running time $\Omega(n^2)$ for every constant number of opinions. A bit surprisingly, our analysis holds even under a stronger model in which, in every round, the activated agent has access to the current opinions of *all* agents in the system.

For comparison, in *symmetric consensus*[3] convergence is achieved in $\mathcal{O}(n \log n)$ rounds with high probability, for a large class of majority-like dynamics and using samples of constant size [Schoenebeck and Yu, 2018]. We thus have an exponential gap between the two settings, in terms of the average number of activations per agent.[4]

We further show that our lower bound is essentially tight. Interestingly, we prove that the standard voter model achieves almost optimal performance, despite using samples of size $\ell = 1$. Specifically, in Section 4, we prove that the voter model converges to the correct opinion within $\mathcal{O}(n^2 \log n)$ rounds in expectation and $\mathcal{O}(n^2 \log^2 n)$ rounds with high probability. This result and the lower bound of Section 3 together suggest that sample size cannot be a key ingredient in achieving fast consensus to the correct opinion after all.

Finally, we argue that allowing agents to use a relatively small amount of memory can drastically decrease convergence time. As mentioned earlier in the introduction, this result has been formally proved in [Korman and Vacus, 2022] in the *parallel* setting, where at every round, all agents are activated simultaneously. We devise a suitable adaptation of the algorithm proposed in [Korman and Vacus, 2022] to work in the sequential, random activation model that is considered in this paper. This adaptation uses samples of size $\ell = \Theta(\log n)$ and $\Theta(\log \log n)$ bits of local memory. Empirical evidence discussed in Section 5 suggests that its convergence time might be compatible with $n \log^{O(1)} n$. In terms of parallel time (i.e., the average number of activations per agent), this would imply an exponential gap between this case and the memoryless case.

1.2 Previous Work

The problem we consider spans a number of areas of potential interest across several communities. The corresponding literature is vast and providing an exhaustive review is infea-

[1] All results we present seamlessly extend (up to constants) to a constant number of source agents.

[2] All the results directly hold also if the sampling is without replacement.

[3] In the remainder, by *symmetric consensus* we mean the standard setting in which agents are required to eventually achieve consensus on *any* of the opinions that are initially present in the system.

[4] This measure is often referred to as the *parallel time* in distributed computing literature [Czumaj and Lingas, 2023].

sible here. In the following paragraphs, we discuss previous contributions that most closely relate to the present work.

Information dissemination with limited communication. Dissemination is especially difficult when communication is limited and/or when the environment is noisy or unpredictable. For this reason, a line of recent work in distributed computing focuses on designing robust protocols, which are tolerant to faults and/or require minimal assumptions on the communication patterns. An effective theoretical framework to address these challenges is that of self-stabilization, in which problems related to the scenario we consider have been investigated, such as self-stabilizing clock synchronization or majority computations [Aspnes and Ruppert, 2009; Ben-Or *et al.*, 2008; Boczkowski *et al.*, 2019]. In general however, these models make few assumptions about memory and/or communication capabilities and they rarely fit the framework of passive communication. Extending the self-stabilization framework to scenarios inspired by biological distributed systems was recently done in [Korman and Vacus, 2022], with interesting preliminary results discussed earlier in the introduction.

Opinion dynamics. Opinion dynamics have been extensively used to investigate opinion formation processes resulting in stable consensus and/or clustering equilibria [Becchetti *et al.*, 2020; Coates *et al.*, 2018; Out and Zehmakan, 2021]. One of the most popular opinion dynamics is the voter model, introduced to study conflicts between species in biology and in interacting particle systems [Clifford and Sudbury, 1973; Holley and Liggett, 1975]. The investigation of majority update rules and variants of it [Becchetti *et al.*, 2020; Berenbrink *et al.*, 2022; Doerr *et al.*, 2011; Mossel *et al.*, 2014] originated from the study of consensus processes in spin systems [Krapivsky and Redner, 2003].

The recent past has witnessed increasing interest for biased variants of opinion dynamics [Anagnostopoulos *et al.*, 2020; Berenbrink *et al.*, 2022; Cruciani *et al.*, 2021; Lesfari *et al.*, 2022]. In general, the focus of this line of work is different from ours, mostly being on the sometimes complex interplay between bias and convergence to an equilibrium, possibly represented by global adoption of one of the opinions. In contrast, our focus is on how quickly dynamics can converge to the (unknown) correct opinion when agents only have access to random samples of the opinions held by other agents.[5]

Consensus in the presence of zealot agents. A large body of work considers opinion dynamics in the presence of zealot agents, i.e., agents (generally holding heterogeneous opinions) that never depart from their initial opinion [D'Amore *et al.*, 2022; Masuda, 2015; Moreno *et al.*, 2020; Yildiz *et al.*, 2013] and may try to influence the rest of the agent population. In this case, the process resulting from a certain dynamics can result in equilibria characterized by multiple opinions. To the best of our knowledge, the main focus of this body of work is different from ours, mostly concerning the impact of the number of zealots, their positions in the network and the topology of the network itself on such equilibria [Fudolig and Esguerra, 2014; Masuda, 2015; Moreno *et al.*, 2020; Yildiz *et al.*, 2013].

[5]For reference, it is easy to verify that majority or best-of-k majority rules [Schoenebeck and Yu, 2018] (which have frequently been considered in the above literature) in general fail to complete the dissemination task we consider.

2 Notations and Preliminaries

We consider a system consisting of n *anonymous* agents. We denote by $x_u^{(t)}$ the opinion held by agent u at the end of round t, dropping the superscript whenever it is clear from the context. The *configuration* of the system at round t is the vector $\mathbf{x^{(t)}}$ with n entries, such that its u-th entry is $x_u^{(t)}$.

We are interested in dynamics that efficiently *disseminate* the correct opinion. I.e., (i) they eventually bring the system into the *correct configuration* in which all agents share the correct opinion, and (ii) they do so in as few rounds as possible. For brevity, we sometimes refer to the latter quantity as *convergence time*. If T is the convergence time of an execution, we denote by T/n the average number of activations per agent, a measure often referred to as *parallel time* in the distributed computing literature [Czumaj and Lingas, 2023]. For ease of exposition, in the remainder we assume that opinions are binary (i.e., they belong to $\{1, 0\}$). We remark the following: (i) the lower bound on the convergence time given in Section 3 already applies by restricting attention to the binary case, and, (ii) it is easy to extend the analysis of the voter model given in Section 4 to the general case of k opinions using standard arguments.

Memoryless dynamics. We consider dynamics in which, beyond being anonymous, non-source agents are memoryless and identical. We capture these and the general requirements outlined in Section 1 by the following decision rule, describing the behavior of agent u

1. u is presented with a uniform sample S of size ℓ;
2. u adopts opinion 1 with probability $g_{x_u}(|S|)$, where $|S|$ denotes the number of 1's in sample S.

Here, $g_{x_u} : \{0, \dots, \ell\} \to [0, 1]$ is a function that assigns a probability value to the number of ones that appear in S. In particular, g_{x_u} assigns probability zero to opinions with no support in the sample, i.e., $g_{x_u}(0) = 0$ and $g_{x_u}(\ell) = 1$.[6] Note that, in principle, g_{x_u} may depend on the current opinion of agent u.

The class of dynamics described by the general rule above strictly includes all memoryless algorithms that are based on random samples of fixed size including popular dynamics, such as the voter model and a large class of quasi-majority dynamics [Liggett, 2012; Schoenebeck and Yu, 2018; Becchetti *et al.*, 2020].

Markov chains. In the remainder, we consider discrete time, discrete space Markov chains, whose state space is represented by an integer interval $\chi = \{z, z + 1, \dots, n\}$, for suitable $z \geqslant 1$ and $n > z$, without loss of generality (the reason for this labeling of the states will be clear in the next sections). Let X_t be the random variable that represents the

[6]In general, dynamics not meeting this constraint cannot enforce consensus.

state of the chain at round $t \geqslant 0$. The *hitting time* [Levin and Peres, 2017, Section 1] of state $x \in S$ is the first time the chain is in state x, namely:

$$\tau_x = \min\{t \geqslant 0 : X_t = x\}.^7$$

A basic ingredient used in this paper is describing the dynamics we consider in terms of suitable *birth-death* chains, in which the only possible transitions from a given state $i \geqslant z$ are to the states $i, i+1$ (if $i \leqslant n-1$) and $i-1$ (if $i \geqslant z+1$). In the remainder, we denote by p_i and q_i respectively the probability of moving to $i+1$ and the probability of moving to $i-1$ when the chain is in state i. Note that $p_n = 0$ and $q_z = 0$. Finally, $r_i = 1 - p_i - q_i$ denotes the probability that, when in state i, the chain remains in that state in the next step.

A birth-death chain for memoryless dynamics. The global behaviour of a system with z source agents holding opinion (wlog) 1 and in which all other agents revise their opinions according to the general dynamics described earlier when activated, is completely described by a birth-death chain $\mathcal{C}_1$ with state space $\{z, ..., n\}$ and the following transition probabilities, for $i = z, \dots n-1$:

$$\begin{aligned} p_i &= \mathbb{P}\left(X_{t+1} = i+1 \mid X_t = i\right) \\ &= \frac{n-i}{n} \sum_{s=0}^{\ell} g_0(s) \mathbb{P}\left(|S| = s \mid X_t = i\right) \\ &= \frac{n-i}{n} \mathbb{E}_i\left[g_0(|S|)\right], \end{aligned} \tag{1}$$

where X_t is simply the number of agents holding opinion 1 at the end of round t and where, following the notation of [Levin and Peres, 2017], for a random variable V defined over some Markov chain $\mathcal{C}$, we denote by $\mathbb{E}_i[V]$ the expectation of V when $\mathcal{C}$ starts in state i. Eq. (1) follows from the law of total probability applied to the possible values for $|S|$ and observing that (a) the transition $i \to i+1$ can only occur if an agent holding opinion 0 is selected for update, which happens with probability $(n-i)/n$, and (b) if such an agent observes s agents with opinion 1 in its sample, it will adopt that opinion with probability $g_0(s)$. Likewise, for $i = z+1, \dots, n-1$:

$$q_i = \mathbb{P}\left(X_{t+1} = i-1 \mid X_t = i\right) = \frac{i-z}{n}(1 - \mathbb{E}_i\left[g_1(|S|)\right]), \tag{2}$$

with the only caveat that, differently from the previous case, the transition $i+1 \to i$ can only occur if an agent with opinion 1 is selected for update and *this agent is not a source*. For this chain, in addition to $p_n = 0$ and $q_z = 0$ we also have $q_n = 0$, which follows since $g_1(\ell) = 1$.

We finally note the following (obvious) connections between $\mathcal{C}_1$ and any specific opinion dynamics P: (i) the specific birth-death chain for P is obtained from $\mathcal{C}_1$ by specifying the corresponding g_0 and g_1 in Eqs. (1) and (2) above; and (ii) the expected convergence time of P starting in a configuration with $i \geqslant z$ agents holding opinion 1 is simply $\mathbb{E}_i[\tau_n]$.

[7] Note that the hitting time in general depends on the initial state. Following [Levin and Peres, 2017], we specify it when needed.

3 Lower Bound

In this section, we prove a lower bound on the convergence time of memoryless dynamics. We show that this negative result holds in a very-strong sense: any dynamics must take $\Omega(n^2)$ rounds in expectation, even if the agents have full knowledge of the current system configuration.

As mentioned in the previous section, we restrict the analysis to the case of two opinions, namely 0 and 1, w.l.o.g. To account for the fact that agents have access to the exact configuration of the system, we slightly modify the notation introduced in Section 2, so that here $g_{x_u} : \{0, \dots, n\} \to [0, 1]$ assigns a probability to the number of ones that appear in the population, rather than in a random sample of size ℓ. Before we prove our main result, we need the following two technical lemmas. Their proofs are omitted due to space constraints but they can be found in [Becchetti *et al.*, 2023].

Lemma 1. *For every $N \in \mathbb{N}$, for every $x \in \mathbb{R}^N$ s.t. for every $i \in \{1, \dots, N\}$, $x_i > 0$, we have either $\sum_{i=1}^{N} x_i \geqslant N$ or $\sum_{i=1}^{N} \frac{1}{x_i} \geqslant N$.*

Lemma 2. *Consider any birth-death chain on $\{0, \dots, n\}$. For $1 \leqslant i \leqslant j \leqslant n$, let $a_i = q_i/p_{i-1}$ and $a(i:j) = \prod_{k=i}^{j} a_k$. Then, $\mathbb{E}_0[\tau_n] \geqslant \sum_{1 \leqslant i < j \leqslant n} a(i:j)$.*

Theorem 3. *Fix $z \in \mathbb{N}$. In the presence of z source agents, the expected convergence time of any memoryless dynamics is at least $\Omega(n^2)$, even when each sample contains the complete configuration of the opinions in the system, i.e., the case $\ell = n$.*

Proof. Fix $z \in \mathbb{N}$. Let $n \in \mathbb{N}$, s.t. $n > 4z$, and let P be any memoryless dynamics. The idea of the proof is to show that the birth-death chain associated with P, obtained from the chain $\mathcal{C}_1$ described in Section 2 by specifying g_0 and g_1 for the dynamics P, cannot be "fast" in both directions at the same time. We restrict the analysis to the subset of states $\chi = \{n/4, \dots, 3n/4\}$. More precisely, we consider the two following birth-death chains: (i) $\mathcal{C}$ with state space χ, whose states represent the number of agents with opinion 1, and assuming that the source agents hold opinion 1; and (ii) $\mathcal{C}'$ with state space χ, whose states represent the number of agents with opinion 0, and assuming that the source agents hold opinion 0. Let $\tau_{3n/4}$ (resp. $\tau'_{3n/4}$) be the hitting time of the state $3n/4$ of chain $\mathcal{C}$ (resp. $\mathcal{C}'$). We will show that

$$\max\left(\mathbb{E}_{n/4}\left[\tau_{3n/4}\right], \mathbb{E}_{n/4}\left[\tau'_{3n/4}\right]\right) = \Omega(n^2).$$

Let $g_0, g_1 : \chi \to [0, 1]$ be the functions describing P over χ. Following Eqs. (1) and (2) in Section 2, we can derive the transition probabilities for $\mathcal{C}$ as

$$p_i = \frac{n-i}{n} g_0(i), \qquad q_i = \frac{i-z}{n}(1 - g_1(i)). \tag{3}$$

Note that the expectations have been removed as a consequence of agents having "full knowledge" of the configuration. Similarly, for $\mathcal{C}'$, the transition probabilities are

$$p'_i = \frac{n-i}{n}(1 - g_1(n-i)), \qquad q'_i = \frac{i-z}{n} g_0(n-i). \tag{4}$$

Following the definition in the statement of Lemma 2, we define a_i and a'_i for $\mathcal{C}$ and $\mathcal{C}'$ respectively. We have

$$a_i = \frac{q_i}{p_{i-1}} = \frac{i-z}{n-i+1} \cdot \frac{1-g_1(i)}{g_0(i-1)},$$

and

$$a'_i = \frac{q'_i}{p'_{i-1}} = \frac{i-z}{n-i+1} \cdot \frac{g_0(n-i)}{1-g_1(n-i+1)}.$$

Observe that we can multiply these quantities by pairs to cancel the factors on the right hand side:

$$a_{n-i+1} \cdot a'_i = \frac{i-z}{i} \cdot \frac{n-i+1-z}{n-i+1}. \tag{5}$$

$(i-z)/i$ is increasing in i, so it is minimized on χ for $i = n/4$. Similarly, $(n-i+1-z)/(n-i+1)$ is minimized for $i = 3n/4$. Hence, we get the following (rough) lower bound from Eq. (5): for every $i \in \chi$,

$$a_{n-i+1} \cdot a'_i \geqslant \left(1 - \frac{4z}{n}\right)^2. \tag{6}$$

Following the definition in the statement of Lemma 2, we define $a(i:j)$ and $a'(i:j)$ for $\mathcal{C}$ and $\mathcal{C}'$ respectively. From Eq. (6), we get for any $i, j \in \chi$ with $i \leqslant j$:

$$\begin{aligned} a'(i:j) &\geqslant \left(1 - \frac{4z}{n}\right)^{2(j-i+1)} \frac{1}{a(n-j+1:n-i+1)} \\ &\geqslant \left(1 - \frac{4z}{n}\right)^{n} \frac{1}{a(n-j+1:n-i+1)}. \end{aligned}$$

Let $c = c(z) = \exp(-4z)/2$. For n large enough,

$$a'(i:j) \geqslant \frac{c}{a(n-j+1:n-i+1)}. \tag{7}$$

Let $N = n^2/8 + n/4$. By Lemma 1, either

$$\sum_{\substack{i,j\in\chi \\ i<j}} a(i:j) \geqslant N,$$

or (by Eq. (7))

$$\sum_{\substack{i,j\in\chi \\ i<j}} a'(i:j) \geqslant c \sum_{\substack{i,j\in\chi \\ i<j}} \frac{1}{a(i:j)} \geqslant cN.$$

By Lemma 2, it implies that either

$$\mathbb{E}_{n/4}\left[\tau_{3n/4}\right] \geqslant N, \quad \text{or} \quad \mathbb{E}_{n/4}\left[\tau'_{3n/4}\right] \geqslant cN.$$

In both cases, there exists an initial configuration for which at least $\Omega(n^2)$ rounds are needed to achieve consensus, which concludes the proof of Theorem 3. □

4 The Voter Model Is (Almost) Optimal

The voter model is the popular dynamics in which the random agent v, activated at round t, pulls another agent $u \in V$ u.a.r. and updates its opinion to the opinion of u.

In this section, we prove that this dynamics achieves consensus within $\mathcal{O}(n^2 \log n)$ rounds in expectation. We prove the result for $z = 1$, noting that the upper bound can only improve for $z > 1$. Without loss of generality, we assume that 1 is the correct opinion.

The modified chain $\mathcal{C}_2$. In principle, we could study convergence of the voter model using the chain $\mathcal{C}_1$ introduced in Section 2. Unfortunately, $\mathcal{C}_1$ has one absorbing state (the state n corresponding to consensus), hence it is not reversible, so that we cannot leverage known properties of reversible birth-death chains [Levin and Peres, 2017, Section 2.5] that would simplify the proof. Note, however, that we are interested in τ_n, the number of rounds to reach state n under the voter model. To this purpose, it is possible to consider a second chain $\mathcal{C}_2$ that is almost identical to $\mathcal{C}_1$ but reversible. In particular, the transition probabilities p_i and q_i of $\mathcal{C}_2$ are the same as in $\mathcal{C}_1$, for $i = z, \ldots, n-1$. Moreover, we have $p_n = 0$ (as in $\mathcal{C}_1$) but $q_n = 1$.[8] Obviously, for any initial state $i \leqslant n-1$, τ_n has exactly the same distribution in $\mathcal{C}_1$ and $\mathcal{C}_2$. For this reason, in the remainder of this section we consider the chain $\mathcal{C}_2$, unless otherwise stated.

Theorem 4. *For $z = 1$, the voter model achieves consensus to opinion 1 within $\mathcal{O}(n^2 \log n)$ rounds in expectation and within $\mathcal{O}\left(n^2 \log n \log \frac{1}{\delta}\right)$ rounds with probability at least $1 - \delta$, for $0 < \delta < 1$.*

Proof. We first compute the general expression for $\mathbb{E}_z[\tau_n]$, i.e., the expected time to reach state n (thus, consensus) in $\mathcal{C}_2$ when the initial state is z, corresponding to the system starting in a state in which only the source agents hold opinion 1. We then give a specific upper bound when $z = 1$. First of all, we recall that, for z source agents we have that $\mathbb{E}_z[\tau_n] = \sum_{k=z+1}^{n} \mathbb{E}_{k-1}[\tau_k]$. Considering the general expressions of the p_i's and q_i's in Eq. (1) and Eq. (2), we soon observe that for the voter model $g_0 = g_1 = g$, since the output does not depend on the opinion of the agent, and $\mathbb{E}[g(|S|)] = i/n$ whenever the number of agent with opinion 1 in the system is i. Hence for $\mathcal{C}_2$ we have

$$p_i = \begin{cases} \frac{(n-i)i}{n^2}, & \text{for } i = z, \ldots, n-1 \\ 0, & \text{for } i = n \end{cases}$$
$$q_i = \begin{cases} 0, & \text{for } i = z \\ \frac{(n-i)(i-z)}{n^2}, & \text{for } i = z+1, \ldots, n-1 \\ 1, & \text{for } i = n. \end{cases} \tag{8}$$

The proof now proceeds along the following steps.

General expression for $\mathbb{E}_{k-1}[\tau_k]$. It is not difficult to see that

$$\mathbb{E}_{k-1}[\tau_k] = \frac{1}{q_k w_k} \sum_{j=z}^{k-1} w_j\,, \tag{9}$$

where $w_0 = 1$ and $w_k = \prod_{i=z+1}^{k} \frac{p_{i-1}}{q_i}$, for $k = z+1, \ldots, n$. Indeed, the w_k's satisfy the detailed balanced conditions $p_{k-1} w_{k-1} = q_k w_k$ for $k = z+1, \ldots, n$, since

$$p_{k-1} w_{k-1} = p_{k-1} \frac{q_k}{p_{k-1}} \prod_{i=z+1}^{k} \frac{p_{i-1}}{q_i} = q_k w_k.$$

and Eq. (9) follows proceeding like in [Levin and Peres, 2017, Section 2.5].

[8] Setting $q_n = 1$ is only for the sake of simplicity, any positive value will do.

Computing $\mathbb{E}_{k-1}[\tau_k]$ for $\mathcal{C}_2$. First of all, considering the expressions of p_i and q_i in Eq. (8), for $k = z+1, \dots, n-1$ we have

$$w_k = \prod_{i=z+1}^{k} \frac{(n-i+1)(i-1)}{(i-z)(n-i)}$$
$$= \prod_{i=z+1}^{k} \frac{n-i+1}{n-i} \cdot \prod_{i=z+1}^{k} \frac{i-1}{i-z} = \frac{n-z}{n-k} \cdot \prod_{i=z+1}^{k} \frac{i-1}{i-z}.$$

Hence

$$w_k = \begin{cases} \frac{n-z}{n-k} f(k), & \text{for } k = z+1, \dots, n-1 \\ \frac{(n-z)(n-1)}{n^2} f(n-1), & \text{for } k = n, \end{cases}$$

where $f(k) = \prod_{i=z+1}^{k} \frac{i-1}{i-z}$.

The case $z = 1$. In this case, the formulas above simplify and, for $k = z+1, \dots, n-1$, we have

$$\mathbb{E}_{k-1}[\tau_k] = \frac{n^2}{(k-1)f(k)} \sum_{j=1}^{k-1} \frac{f(j)}{n-j} = \frac{n^2}{k-1} \sum_{j=1}^{k-1} \frac{1}{n-j},$$

where the last equality follows from the fact that $f(z) = f(z+1) = \cdots = f(k) = 1$, whenever $z = 1$. Moreover, for $k = n$ we have

$$\mathbb{E}_{n-1}[\tau_n] = \frac{1}{q_n w_n} \sum_{j=1}^{n-1} w_j = \left(\frac{n}{n-1}\right)^2 \sum_{j=1}^{n-1} \frac{n-1}{n-j}$$
$$= \frac{n}{n-1} H_{n-1} = \mathcal{O}(\log n),$$

where H_k denotes the k-th harmonic number. Hence, for $z = 1$ we have

$$\mathbb{E}_1[\tau_n] = \sum_{k=2}^{n} \mathbb{E}_{k-1}[\tau_k]$$
$$= n^2 \sum_{k=2}^{n-1} \frac{1}{k-1} \sum_{j=1}^{k-1} \frac{1}{n-j} + \mathcal{O}(\log n), \quad (10)$$

where in the second equality we took into account that $\mathbb{E}_{n-1}[\tau_n] = \mathcal{O}(\log n)$. Finally, it is easy to see that

$$\sum_{k=2}^{n-1} \frac{1}{k-1} \sum_{j=1}^{k-1} \frac{1}{n-j} = \mathcal{O}(\log n) \quad (11)$$

Indeed, if we split the sum at $\lfloor n/2 \rfloor$, for $k \leqslant \lfloor n/2 \rfloor$ we have

$$\sum_{k=2}^{\lfloor n/2 \rfloor} \frac{1}{k-1} \sum_{j=1}^{k-1} \frac{1}{n-j} \leqslant \sum_{k=2}^{\lfloor n/2 \rfloor} \frac{1}{k-1} \sum_{j=1}^{k-1} \frac{2}{n} = \mathcal{O}(1) \quad (12)$$

and for $k > \lfloor n/2 \rfloor$ we have

$$\sum_{k=\lfloor n/2 \rfloor+1}^{n-1} \frac{1}{k-1} \sum_{j=1}^{k-1} \frac{1}{n-j} \leqslant \sum_{k=\lfloor n/2 \rfloor+1}^{n-1} \frac{2}{n} \sum_{j=0}^{n-1} \frac{1}{n-j}$$
$$= \sum_{k=\lfloor n/2 \rfloor+1}^{n-1} \frac{2}{n} H_n = \mathcal{O}(\log n). \quad (13)$$

From Eqs. (12) and (13) we get Eq. (11), and the first part of the claim follows by using in Eq. (10) the bound in Eq. (11).

To prove the second part of the claim, we use a standard argument. Consider $\lceil \log \frac{1}{\delta} \rceil$ consecutive time intervals, each consisting of $s = 2\lceil \mathbb{E}_1[\tau_n] \rceil = \mathcal{O}(n^2 \log n)$ consecutive rounds. For $i = 1, \dots, s-1$, if the chain did not reach state n in any of the first $i-1$ intervals, then the probability that the chain does not reach state n in the i-th interval is at most $1/2$ by Markov's inequality. Hence, the probability that the chain does not reach state n in any of the intervals is at most $(1/2)^{\log(1/\delta)} = \delta$. □

4.1 Handling Multiple Opinions

Consider the case in which the set of possible opinions is $\{1, \dots, k\}$ for $k \geqslant 2$, with 1 the correct opinion without loss of generality. We collapse opinions $2, \dots, k$ into one class, i.e., opinion 0 without loss of generality. We then consider the random variable X_t, giving the number of agents holding opinion 1 at the end of round t. Clearly, the configuration in which all agents hold opinion 1 is the only absorbing state under the voter model and convergence time is defined as $\min\{t \geqslant 0 : X_t = n\}$. For a generic number i of agents holding opinion 1, we next compute the probability p_i of the transition $i \to i+1$ (for $i \leqslant n-1$) and the probability q_i of the transition $i \to i-1$ (for $i \geqslant z+1$):

$$p_i = \mathbb{P}(X_{t+1} = i+1 \mid X_t = i) = \frac{n-i}{n} \cdot \frac{i}{n},$$

where the first factor in the right hand side of the above equality is the probability of activating an agent holding an opinion other than 1, while the second factor is the probability that said agent in turn copies the opinion of an agent holding the correct opinion. Similarly, we have:

$$q_i = \mathbb{P}(X_{t+1} = i-1 \mid X_t = i) = \frac{i-z}{n} \cdot \frac{n-i}{n},$$

with the first factor in the right hand side the probability of sampling a non-source agent holding opinion 1 and the second factor the probability of this agent in turn copying the opinion of an agent holding any opinions other than 1.

The above argument implies that if we are interested in the time to converge to the correct opinion, variable X_t is what we are actually interested in. On the other hand, it is immediately clear that the evolution of X_t is described by the birth-death chain $\mathcal{C}_1$ introduced in Section 2 (again with n as the only absorbing state) or by its reversible counterpart $\mathcal{C}_2$. This in turn implies that the analysis of Section 4 seamlessly carries over to the case of multiple opinions.

5 Faster Dissemination With Memory

In this section, we give experimental evidence suggesting that dynamics using a modest amount of memory can achieve consensus in an almost linear number of rounds.

The dynamics that we use is derived from the algorithm of [Korman and Vacus, 2022], and called "Follow The Trend" (FtT). It uses a sample size of $\ell = 10 \log n$ and works for an

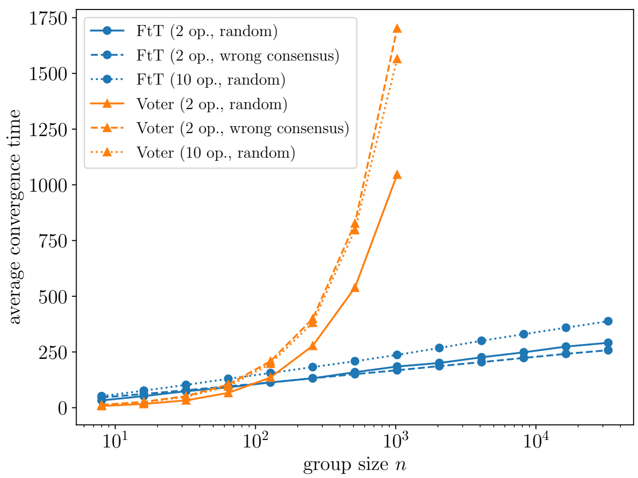

Figure 1: **"Follow the Trend" versus the voter model.** *Average convergence time (in parallel rounds) is depicted for different values of* n*, over 100 iterations each, for* $z = 1$ *source agent. Blue lines with circular markers correspond to our candidate dynamics (FtT). Orange lines with triangular markers correspond to the voter model. Full lines depict initial configurations in which all opinions are drawn uniformly at random from* $\{0, 1\}$*. Dashed lines depict initial configurations in which the correct opinion is* 0 *and all other opinions are* 1*. Dotted lines depict initial configurations in which all opinions are drawn uniformly at random from* $\{0, \ldots, 9\}$*.*

arbitrary number of opinions k^9. We start by giving an informal description of the dynamics. While each agent is activated once every n rounds in average according to the model definition, it only becomes *busy* once every ℓ activations when following FtT (this behavior is implemented by an internal clock). Let us denote by $A(s, i, u)$ the number of samples corresponding to opinion i observed by Agent u during its s^{th} "busy" activation. Upon the s^{th} busy activation, Agent u adopts the opinion maximizing $A(s, i, u) - A(s-1, i, u)$ (hence following the "emerging trend"). In case of a tie, two scenarios are possible: if the current opinion of the agent maximizes $A(s, i, u) - A(s-1, i, u)$, then the agent remains with it; otherwise, the tie is broken uniformly at random.

The FtT dynamics is then compared experimentally to the voter model. Results are summed up in Figure 1, in terms of parallel rounds (one parallel round corresponds to n activations). They suggest that the expected convergence time of FtT is about $\Theta(\text{polylog} n)$ parallel rounds. In terms of parallel time, this represents an exponential gap when compared to the lower bound in Theorem 3 established for memoryless dynamics.

More details regarding the dynamics and its simulations can be found in [Becchetti *et al.*, 2023].

[9] The factor 10 in the sample size can be replaced by any sufficiently large constant.

6 Discussion and Future Work

This work investigates the role played by memory in multi-agent systems that rely on passive communication and aim to achieve consensus on an opinion held by few "knowledgable" individuals [Korman and Vacus, 2022; Couzin *et al.*, 2005; Ayalon *et al.*, 2021]. Under the model we consider, we prove that incorporating past observations in the current decision is necessary for achieving fast convergence even if the observations regarding the current opinion configuration are complete. The same lower bound proof can in fact be adapted to any process that is required to alternate the consensus (or semi-consensus) opinion, i.e., to let the population agree (or almost agree) on one opinion, and then let it agree on the other opinion, and so forth. Such oscillating behaviour is fundamental to sequential decision making processes [Ayalon *et al.*, 2021].

The ultimate goal of this line of research is to reflect on biological processes and derive lower bounds on biological parameters. However, despite the generality of our model, more work must be done to obtain concrete biological conclusions. Conducting an experiment that fully adheres to our model, or refining our results to apply to more realistic settings remains for future work. Candidate experimental settings that appear to be promising include fish schooling [Couzin *et al.*, 2005; Couzin *et al.*, 2011],collective sequential decision-making in ants [Ayalon *et al.*, 2021], and recruitment in ants [Razin *et al.*, 2013]. If successful, such an outcome would be highly pioneering from a methodological perspective. Indeed, to the best of our knowledge, a concrete lower bound on a biological parameter that is achieved in an indirect manner via mathematical considerations has never been obtained.

Acknowledgements

This project has received funding from:

- The European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 834861).
- The European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 648032).
- Spoke 1 "FutureHPC & BigData" of the *Italian Research Center on High-Performance Computing, Big Data and Quantum Computing (ICSC)* funded by *MUR Missione 4 Componente 2 Investimento 1.4: Potenziamento strutture di ricerca e creazione di "campioni nazionali" di R&S (M4C2-19) - Next Generation EU (NGEU).*
- Partially supported by the ERC Advanced Grant 788893 AMDROMA "Algorithmic and Mechanism Design Research in Online Markets", the EC H2020RIA project "SoBigData++" (871042), and the MIUR PRIN project ALGADIMAR "Algorithms, Games, and Digital Markets.

References

[Anagnostopoulos *et al.*, 2020] Aris Anagnostopoulos, Luca Becchetti, Emilio Cruciani, Francesco Pasquale, and Sara

Rizzo. Biased opinion dynamics: When the devil is in the details. In *29th International Joint Conference on Artificial Intelligence*, pages 53–59, 2020.

[Aspnes and Ruppert, 2009] James Aspnes and Eric Ruppert. An introduction to population protocols. *Middleware for Network Eccentric and Mobile Applications*, pages 97–120, 2009.

[Ayalon *et al.*, 2021] Oran Ayalon, Yigal Sternklar, Ehud Fonio, Amos Korman, Nir S Gov, and Ofer Feinerman. Sequential decision-making in ants and implications to the evidence accumulation decision model. *Frontiers in applied mathematics and statistics*, page 37, 2021.

[Bastide *et al.*, 2021] Paul Bastide, George Giakkoupis, and Hayk Saribekyan. Self-stabilizing clock synchronization with 1-bit messages. In *ACM-SIAM Symposium on Discrete Algorithms*, pages 2154–2173. SIAM, 2021.

[Becchetti *et al.*, 2020] Luca Becchetti, Andrea E. F. Clementi, and Emanuele Natale. Consensus dynamics: An overview. *SIGACT News*, 51(1):58–104, 2020.

[Becchetti *et al.*, 2023] Luca Becchetti, Andrea Clementi, Amos Korman, Francesco Pasquale, Luca Trevisan, and Robin Vacus. On the role of memory in robust opinion dynamics. *arXiv preprint arXiv:2302.08600*, 2023.

[Ben-Or *et al.*, 2008] Michael Ben-Or, Danny Dolev, and Ezra N Hoch. Fast self-stabilizing byzantine tolerant digital clock synchronization. In *27th ACM symposium on Principles of Distributed Computing*, pages 385–394, 2008.

[Berenbrink *et al.*, 2022] Petra Berenbrink, Martin Hoefer, Dominik Kaaser, Pascal Lenzner, Malin Rau, and Daniel Schmand. Asynchronous opinion dynamics in social networks. In *21st International Conference on Autonomous Agents and Multiagent Systems*, pages 109–117, 2022.

[Bialek *et al.*, 2012] William Bialek, Andrea Cavagna, Irene Giardina, Thierry Mora, Edmondo Silvestri, Massimiliano Viale, and Aleksandra M Walczak. Statistical mechanics for natural flocks of birds. *Proceedings of the National Academy of Sciences*, 109(13):4786–4791, 2012.

[Boczkowski *et al.*, 2018] Lucas Boczkowski, Emanuele Natale, Ofer Feinerman, and Amos Korman. Limits on reliable information flows through stochastic populations. *PLoS computational biology*, 14(6):e1006195, 2018.

[Boczkowski *et al.*, 2019] Lucas Boczkowski, Amos Korman, and Emanuele Natale. Minimizing message size in stochastic communication patterns: fast self-stabilizing protocols with 3 bits. *Distributed Computing*, 32(3), 2019.

[Clifford and Sudbury, 1973] Peter Clifford and Aidan Sudbury. A model for spatial conflict. *Biometrika*, 60(3):581–588, 1973.

[Coates *et al.*, 2018] Adam Coates, Liangxiu Han, and Anthony Kleerekoper. A unified framework for opinion dynamics. In *17th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS 2018*, pages 1079–1086, 2018.

[Couzin *et al.*, 2005] Iain D Couzin, Jens Krause, Nigel R Franks, and Simon A Levin. Effective leadership and decision-making in animal groups on the move. *Nature*, 433(7025):513–516, 2005.

[Couzin *et al.*, 2011] Iain D Couzin, Christos C Ioannou, Güven Demirel, Thilo Gross, Colin J Torney, Andrew Hartnett, Larissa Conradt, Simon A Levin, and Naomi E Leonard. Uninformed individuals promote democratic consensus in animal groups. *Science*, 334(6062):1578–1580, 2011.

[Cruciani *et al.*, 2021] Emilio Cruciani, Hlafo Alfie Mimun, Matteo Quattropani, and Sara Rizzo. Phase transitions of the k-majority dynamics in a biased communication model. In *ACM International Conference on Distributed Computing and Networking*, pages 146–155. ACM, 2021.

[Cvikel *et al.*, 2015] Noam Cvikel, Katya Egert Berg, Eran Levin, Edward Hurme, Ivailo Borissov, Arjan Boonman, Eran Amichai, and Yossi Yovel. Bats aggregate to improve prey search but might be impaired when their density becomes too high. *Current Biology*, 25(2):206–211, 2015.

[Czumaj and Lingas, 2023] Artur Czumaj and Andrzej Lingas. On parallel time in population protocols. *Information Processing Letters*, 179:106314, 2023.

[D'Amore *et al.*, 2022] Francesco D'Amore, Andrea E. F. Clementi, and Emanuele Natale. Phase transition of a nonlinear opinion dynamics with noisy interactions. *Swarm Intell.*, 16(4):261–304, 2022.

[Doerr *et al.*, 2011] Benjamin Doerr, Leslie Ann Goldberg, Lorenz Minder, Thomas Sauerwald, and Christian Scheideler. Stabilizing consensus with the power of two choices. In *23rd ACM Symposium on Parallelism in Algorithms and Architectures*, 2011.

[Feinerman and Korman, 2017a] Ofer Feinerman and Amos Korman. The ants problem. *Distributed Computing*, 30(3):149–168, 2017.

[Feinerman and Korman, 2017b] Ofer Feinerman and Amos Korman. Individual versus collective cognition in social insects. *Journal of Experimental Biology*, 220(1):73–82, 2017.

[Fonio *et al.*, 2016] Ehud Fonio, Yael Heyman, Lucas Boczkowski, Aviram Gelblum, Adrian Kosowski, Amos Korman, and Ofer Feinerman. A locally-blazed ant trail achieves efficient collective navigation despite limited information. *Elife*, 5:e20185, 2016.

[Franks *et al.*, 2002] Nigel R Franks, Stephen C Pratt, Eamonn B Mallon, Nicholas F Britton, and David JT Sumpter. Information flow, opinion polling and collective intelligence in house–hunting social insects. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 357(1427):1567–1583, 2002.

[Fudolig and Esguerra, 2014] Mikaela Irene D Fudolig and Jose Perico H Esguerra. Analytic treatment of consensus achievement in the single-type zealotry voter model. *Physica A: Statistical Mechanics and its Applications*, 413:626–634, 2014.

[Gelblum *et al.*, 2015] Aviram Gelblum, Itai Pinkoviezky, Ehud Fonio, Abhijit Ghosh, Nir Gov, and Ofer Feinerman. Ant groups optimally amplify the effect of transiently informed individuals. *Nature Communications*, 6(1):1–9, 2015.

[Giraldeau and Caraco, 2018] Luc-Alain Giraldeau and Thomas Caraco. Social foraging theory. In *Social Foraging Theory*. Princeton University Press, 2018.

[Guinard and Korman, 2021] Brieuc Guinard and Amos Korman. Intermittent inverse-square lévy walks are optimal for finding targets of all sizes. *Science Advances*, 7(15):eabe8211, 2021.

[Holley and Liggett, 1975] Richard A Holley and Thomas M Liggett. Ergodic theorems for weakly interacting infinite systems and the voter model. *The Annals of Probability*, pages 643–663, 1975.

[Korman and Vacus, 2022] Amos Korman and Robin Vacus. Early adapting to trends: Self-stabilizing information spread using passive communication. In *ACM Symposium on Principles of Distributed Computing*, pages 235–245. ACM, 2022.

[Krapivsky and Redner, 2003] P. L. Krapivsky and S. Redner. Dynamics of majority rule in two-state interacting spin systems. *Phys. Rev. Lett.*, 90:238701, 2003.

[Lesfari *et al.*, 2022] Hicham Lesfari, Frédéric Giroire, and Stéphane Pérennes. Biased majority opinion dynamics: Exploiting graph k-domination. In *31st International Joint Conference on Artificial Intelligence*, pages 377–383. ijcai.org, 2022.

[Levin and Peres, 2017] David A Levin and Yuval Peres. *Markov chains and mixing times*, volume 107. American Mathematical Soc., 2017.

[Liggett, 2012] Thomas Milton Liggett. *Interacting particle systems*, volume 276. Springer Science & Business Media, 2012.

[Lindauer, 1957] M Lindauer. Communication in swarm-bees searching for a new home. *Nature*, 179(4550):63–66, 1957.

[Masuda, 2015] Naoki Masuda. Opinion control in complex networks. *New Journal of Physics*, 17, 03 2015.

[Moreno *et al.*, 2020] Guillermo Romero Moreno, Long Tran-Thanh, and Markus Brede. Continuous influence maximisation for the voter dynamics: is targeting high-degree nodes a good strategy? In *International Conference on Autonomous Agents and Multi-Agent Systems*, 2020.

[Mossel *et al.*, 2014] Elchanan Mossel, Joe Neeman, and Omer Tamuz. Majority dynamics and aggregation of information in social networks. *Auton. Agents Multi Agent Syst.*, 28(3):408–429, 2014.

[Out and Zehmakan, 2021] Charlotte Out and Ahad N. Zehmakan. Majority vote in social networks: Make random friends or be stubborn to overpower elites. In *30th International Joint Conference on Artificial Intelligence*, pages 349–355, 2021.

[Razin *et al.*, 2013] Nitzan Razin, Jean-Pierre Eckmann, and Ofer Feinerman. Desert ants achieve reliable recruitment across noisy interactions. *Journal of the Royal Society Interface*, 10(82):20130079, 2013.

[Schoenebeck and Yu, 2018] Grant Schoenebeck and Fang-Yi Yu. Consensus of interacting particle systems on erdös-rényi graphs. In *Proceedings of the Twenty-Ninth ACM-SIAM Symposium on Discrete Algorithms*, pages 1945–1964, 2018.

[Seeley, 2003] Thomas D Seeley. Consensus building during nest-site selection in honey bee swarms: the expiration of dissent. *Behavioral Ecology and Sociobiology*, 53(6):417–424, 2003.

[Sumpter, 2010] David JT Sumpter. Collective animal behavior. In *Collective Animal Behavior*. Princeton University Press, 2010.

[Wilkinson, 1992] Gerald S Wilkinson. Information transfer at evening bat colonies. *Animal Behaviour*, 44:501–518, 1992.

[Yildiz *et al.*, 2013] Ercan Yildiz, Asuman Ozdaglar, Daron Acemoglu, Amin Saberi, and Anna Scaglione. Binary opinion dynamics with stubborn agents. *ACM Trans. Econ. Comput.*, 1(4), dec 2013.