

On a Voter Model with Context-Dependent Opinion Adoption

Luca Becchetti¹, Vincenzo Bonifaci², Emilio Cruciani³, Francesco Pasquale⁴

¹Sapienza University of Rome – Rome, Italy

²Roma Tre University – Rome, Italy

³Paris-Lodron University of Salzburg – Salzburg, Austria

⁴Tor Vergata University – Rome, Italy

becchetti@diag.uniroma1.it, vincenzo.bonifaci@uniroma3.it,
emilio.cruciani@plus.ac.at, pasquale@mat.uniroma2.it

Abstract

Opinion diffusion is a crucial phenomenon in social networks, often underlying the way in which a collection of agents develops a consensus on relevant decisions. Voter models are well-known theoretical models to study opinion spreading in social networks and structured populations. Their simplest version assumes that an updating agent will adopt the opinion of a neighboring agent chosen at random. These models allow us to study, for example, the probability that a certain opinion will fixate into a consensus opinion, as well as the expected time it takes for a consensus opinion to emerge.

Standard voter models are oblivious to the opinions held by the agents involved in the opinion adoption process. We propose and study a context-dependent opinion spreading process on an arbitrary social graph, in which the probability that an agent abandons opinion a in favor of opinion b depends on both a and b . We discuss the relations of the model with existing voter models and then derive theoretical results for both the fixation probability and the expected consensus time for two opinions, for both the synchronous and the asynchronous update models.

1 Introduction

The *voter model* is a well-studied stochastic process defined on a graph to model the spread of opinions (or genetic mutations, beliefs, practices, etc.) in a population [Liggett, 1985; Hassin and Peleg, 2001]. In a voter model, each node maintains a state, and when a node requires updating, it will import its state from a randomly chosen neighbor. Updates can be asynchronous, with one node activating per step [Liggett, 1985], or synchronous, with all nodes activating in parallel [Hassin and Peleg, 2001]. The voter model on a graph has been introduced in the 1970s to model opinion dynamics. The case of a complete graph is also very well-known in population genetics, where it was introduced even earlier to study the spread of mutations in a population [Ewens, 2012; Nowak, 2006].

Mathematically, among the main quantities of interest in the study of voter models, there are the *fixation probability*

of an opinion—the probability of reaching a configuration in which each node adopts such opinion—and the expected *consensus* (or *absorption*) *time*—the expected number of steps before all nodes agree on an opinion. Such quantities could in principle be computed for any n -node graph by defining a Markov chain on a set of C^n configurations, where C is the number of opinions, but such an approach is computationally infeasible even for moderate values of n . Therefore, a theoretical analysis of a voter process will often focus on obtaining upper and lower bounds for these quantities, still drawing heavily on the theory of Markov chains [Aldous and Fill, 2002; Levin *et al.*, 2009], but with somewhat different approaches and tools for the synchronous and asynchronous cases.

A limitation of the standard voter process is that the dynamics is oblivious to the states of both the agent u that is updating and of the neighbor that u copies its state from, and the copying always occurs. One could easily imagine a situation (for example, in politics) where an agent holding opinion a is more willing to adopt the opinion b of a neighbor rather than to adopt opinion c . In general, the probability of abandoning opinion a in favor of opinion b might depend on both a and b . This motivates the study of *biased* voter models [Berenbrink *et al.*, 2016; Sood *et al.*, 2008] and in particular motivates us to introduce a voter model with an opinion adoption probability that depends on the *context*, that is, on the opinions of *both* agents involved in an opinion spreading step.

We define and study extensions of the voter model that allow the opinion adoption probability to depend on the pair of opinions involved in an update step. We consider both an asynchronous variant and a synchronous variant of a context-dependent voter model with two opinions, 0 and 1. We assume that an agent holding opinion $c \in \{0, 1\}$ is willing to copy the opinion of an agent holding opinion $c' \in \{0, 1\}$ with some probability $\alpha_{c,c'}$, which models the bias in the update. We study both the fixation probabilities and the expected consensus time.

1.1 Our Findings

In general, a seemingly minor feature as the form of bias we consider has a profound impact on the analytical tractability of the resulting model. While the unbiased case¹ can still be

¹The model is unbiased when $\alpha_{0,1} = \alpha_{1,0}$.

connected to a variant of the voter model and analyzed accordingly with some extra work, the same is not possible for the biased case. Specifically, in Section 3 we prove that a lazy variant of the voter model is equivalent (i.e., it produces the same distribution over possible system’s configurations) to the unbiased variant of the model we consider. The proof, given in the full paper version² for the synchronous case, uses a coupling between the Markov chains that describe the two models. For the asynchronous case, this allows us to directly leverage known connections between the asynchronous voter model and random walks (Proposition 1). For the case of the clique, this general result can be improved, providing explicit, tight bounds on expected consensus time (Theorem 1). In the synchronous case, the above connection is not immediate (i.e., Proposition 1 does not apply) and analyzing expected time to consensus requires adapting arguments that have been used for continuous Markov chains to the synchronous, discrete setting (Theorem 2).

The biased case is considerably harder to analyze, the main reason being that it is no longer possible to collapse the Markov chain describing the system (whose state space is in general the exponentially large set of all possible configurations) to a “simpler” chain, e.g., a random walk on the underlying network, not even in the case of the clique. Despite these challenges, some trends emerge from specific cases. Interestingly, it is possible to derive the exact fixation probabilities for the class of regular networks, highlighting a non-linear dependence on the bias (Theorem 3), while an asymptotically tight analysis for the clique (Theorem 4) suggests that the presence of a bias may have a positive impact on achieving faster consensus in dense networks. Though seemingly intuitive, this last aspect is not a shared property of biased opinion models in general [Montanari and Saberi, 2010; Anagnostopoulos *et al.*, 2022]. The behavior of the model is considerably more complex and technically challenging in the synchronous, biased case. In particular, the preliminary results that we obtain highlight a general dependence of fixation probabilities (Proposition 2) and, notably, expected consensus time (Theorem 5) on both the bias and the initial configuration.

1.2 Related Work

For the sake of space, we mostly discuss results that are most closely related to the setting we consider.

Voter and Voter-Like Models. Due to its versatility, the voter model has been defined multiple times across different disciplines and has a vast literature. As mentioned in the introduction, the special case of the voter model on a complete graph was first introduced in mathematical genetics, being closely related to the so-called Wright-Fisher and Moran processes [Moran, 1958; Kimura, 1962; Ewens, 2012]. The first asynchronous formulation of a voter model on a connected graph has been proposed in the probability and statistics community in the 1970s [Liggett, 1985; Donnelly and Welsh, 1983], while [Hassin and Peleg, 2001] was the first study of this model in a synchronous setting. A classic result of these papers is that the fixation probability of

an opinion c is equal to the weighted fraction of nodes holding opinion c , where the weight of a node is given by its degree [Hassin and Peleg, 2001; Sood *et al.*, 2008].

The expected consensus time of the voter model is much more challenging to derive exactly, even for highly structured graphs. In the asynchronous case, it has often been studied by approximating the process with a continuous diffusion partial differential equation [Ewens, 2012; Sood *et al.*, 2008; Baxter *et al.*, 2008]. For two opinions on the complete graph, this yields the approximation

$$T(n) \approx n^2 h(k/n) \quad (1)$$

where $T(n)$ is the expected consensus time on the n -clique, k is the number of nodes initially holding the first opinion and $h(p) = -p \ln p - (1-p) \ln(1-p)$ [Sood *et al.*, 2008; Ewens, 2012]. To the best of our knowledge, however, no error bound was known for such diffusion approximations³. Another approach is to use the duality between voter model and coalescing random walks [Liggett, 1985; Donnelly and Welsh, 1983], which involves no approximations, but the resulting formulas are hard to interpret, and to paraphrase Donnelly and Welsh [Donnelly and Welsh, 1983], “an exact evaluation of the expected absorption time for a general regular graph is a horrendous computation”. As for approximations, expected consensus time for the voter model can be bounded by $\mathcal{O}((d_{\text{avg}}/d_{\text{min}})(n/\Phi))$, where d_{avg} and d_{min} are, respectively, the average and minimum degrees [Berenbrink *et al.*, 2016]. Most relevant to our discussion is the biased voter model considered by Berenbrink *et al.* [Berenbrink *et al.*, 2016], in which the probability of adoption of an alternative opinion c' depends on c' (and only on c'). While our model is different if we consider more than 2 opinions, there are several other differences with respect to [Berenbrink *et al.*, 2016] even in the binary case. In particular, Berenbrink *et al.* only consider the synchronous setting, they assume there is a “preferred” opinion that is never rejected and that there is a constant gap between the adoption probabilities of the preferred opinion and of the non-preferred one. Finally, they only consider the case where the number k of nodes initially holding the preferred opinion is at least $\Omega(\log n)$. Thus, for example, their results do not apply in the neutral case ($\alpha_{01} = \alpha_{10}$) or when k is, say, $\sqrt{\log n}$. Our results for the biased, synchronous case (Theorem 5) are complementary to those in [Berenbrink *et al.*, 2016]. While their results are stronger when the above assumptions hold, ours address the general and challenging case of an arbitrary initial configuration.

Pull vs Push. We only reviewed here models where nodes “pull” the opinion from their neighbor, since both the standard voter model and our generalization follow this rule, but we remark that “push” models, also known as invasion processes, have also been defined and studied on connected graphs [Lieberman *et al.*, 2005; Díaz *et al.*, 2016]. The asynchronous push model is sometimes called the (generalized) Moran process [Díaz *et al.*, 2016; Nowak, 2006]. We remark that while the pull and push models are interchangeable

²<http://arxiv.org/abs/2305.07377>

³For the n -clique, we show that (1) is correct within an additive $\mathcal{O}(n)$ term. See Theorem 1 (Section 3.1).

on regular graphs, on irregular graphs their behavior can be markedly different.

Other Biased Opinion Dynamics. We are aware of only a few analytically rigorous studies of biased opinion dynamics, including biased variants of the voter model [Sood *et al.*, 2008; Berenbrink *et al.*, 2016; Anagnostopoulos *et al.*, 2022; Cruciani *et al.*, 2021; Durocher *et al.*, 2022], sometimes framed within an evolutionary game setting [Montanari and Saberi, 2010]. In general, these contributions address different models, be it because of the way in which bias is incorporated within the voting rule, the opinion dynamics itself or the temporal evolution of the process (e.g., synchronous vs asynchronous). We remark that all these aspects can deeply affect the overall behavior of the resulting dynamics. Specifically, as observed in a number of more or less recent contributions [Anagnostopoulos *et al.*, 2022; Cruciani *et al.*, 2021; Cooper *et al.*, 2018; Hindersin and Traulsen, 2014], even minor changes in the model that would intuitively produce consistent results with a given baseline can actually induce fundamental differences in the overall behavior, so that it is in general hard to predict if and when results for one model more or less straightforwardly carry over to another model, even qualitatively. Less related to the spirit of this work, a large body of research addresses biased opinion dynamics using different approaches, based on approximations and/or numerical simulations. Examples include numerical simulations for large and more complex scenarios [Hindersin and Traulsen, 2014], mean-field or higher-order [Peralta *et al.*, 2021] and/or continuous approximations [Assaf and Mobilia, 2012]. While these approaches can afford investigation of richer and more complex evolutionary game settings (e.g., [Peralta *et al.*, 2021]), they typically require strong simplifying assumptions to ensure tractability, so that it is harder (if not impossible) to derive rigorous results.

2 Model Formulation

Notation. For a natural number k , let $[k] := \{0, 1, 2, \dots, k - 1\}$. If $G = (V, E)$ is a graph, we write $N_G(u)$ (or simply $N(u)$ if G can be inferred from the context) for the set of neighbors of node u in G . We write d_u for the degree of node u .

Model. We define an opinion dynamics model on networks. The parameters of the model are: i) an underlying *topology*, given by a graph G on n nodes, with symmetric adjacency matrix $A = (a_{uv})_{u,v \in [n]}$, where $a_{uv} \in \{0, 1\}$; ii) a number of *opinions* (or *colors*) $C \geq 2$; iii) an *opinion acceptance matrix* $(\alpha_{c,c'})_{c,c' \in [C]}$, where $\alpha_{c,c'} \in [0, 1]$. The initial opinion of each agent (node) u is encoded by some $x_u^{(0)} \in [C]$.

For any node $u \in [n]$, we define an update process $\text{Update}(u)$ consisting of the following steps (summarized in Algorithm 1):

1. **Sample:** Sample a neighbor v of u uniformly at random, i.e., according to the distribution $(a_{u1}/d_u, \dots, a_{un}/d_u)$ where $a_{uv} = 1$ if $\{u, v\} \in E$ and $a_{uv} = 0$ otherwise. Here $d_u = |N(u)| = \sum_{v \in [n]} a_{uv}$ is the degree of u .
2. **Compare:** Compare u 's opinion $c = x_u$ with v 's opinion $c' = x_v$.

Algorithm 1 $\text{Update}(u)$

- 1: **Sample** $v \in N(u)$
 - 2: $c \leftarrow x_u; c' \leftarrow x_v$
 - 3: **Sample** $\theta \in [0, 1]$
 - 4: **if** $\theta < \alpha_{c,c'}$ **then**
 - 5: $x_u \leftarrow x_v$
 - 6: **return** accept
 - 7: **return** reject
-

3. **Accept/reject:** With probability $\alpha_{c,c'}$, set $x_u \leftarrow x_v$; in this case we say u *accepts* v 's opinion. Otherwise, we say u *rejects* v 's opinion.

We consider two variants of the model, differing in how the updates are scheduled. In one iteration of the *asynchronous* variant, $u \in [n]$ is sampled uniformly at random and $\text{Update}(u)$ is applied. In one iteration of the *synchronous* variant, each node $u \in [n]$ applies $\text{Update}(u)$ in parallel. We denote by $x_u^{(t)}$ the random variable encoding the opinion of node u after t iterations of either the synchronous or the asynchronous dynamics (depending on the context).

The acceptance probabilities $\alpha_{c,c'}$ are parameters of the model. We note that the parameters $\alpha_{c,c'}$ with $c = c'$ are irrelevant for the dynamics, since a node sampling a neighbor of identical opinion will not change opinion, irrespective of whether it accepts the neighbor's opinion or not. Hence, to specify the opinion acceptance matrix $C(C - 1)$ parameters are sufficient; we can assume that the diagonal entries are equal to 1. In particular, when $C = 2$ it is enough to specify α_{01} and α_{10} . When $\alpha_{01} = \alpha_{10} = 1$, the model boils down to the standard voter model [Hassin and Peleg, 2001; Liggett, 1985].

In the rest of this work we assume $C = 2$. In this case, we say that the model is *unbiased* if the opinion acceptance matrix is symmetric, i.e., $\alpha_{01} = \alpha_{10}$, and *biased* otherwise.

Quantities of Interest. The *fixation probability* of opinion 1 is the probability that there exists an iteration t such that $x_u^{(t)} = 1$ for all $u \in [n]$. The *consensus time* is the index of the first iteration t such that $x_u^{(t)} = x_v^{(t)}$ for all $u, v \in [n]$.

3 The Unbiased Setting

Before embarking on the biased case, which is substantially more complex, in this section we review or prove directly results for the unbiased setting ($\alpha_{01} = \alpha_{10}$). We consider the asynchronous and the synchronous variants separately. In all formulas of this section, $\alpha = \alpha_{01} = \alpha_{10}$.

3.1 Asynchronous Variant

The main, intuitive observation about the unbiased asynchronous variant of our model is that the model can equivalently be described by a suitable, "lazy" voter model, where each iteration is either an idle iteration (with probability $1 - \alpha$) or an iteration of the standard asynchronous voter model (with probability α).

This in turn implies that, for the fixation probability one can simply disregard the idle iterations and therefore obtain the same fixation probability as for the standard asynchronous

voter model. In an arbitrary topology, this was derived by Sood et al. [Sood et al., 2008]: if we call ϕ^{avoter} the fixation probability of the asynchronous voter model, then

$$\phi^{\text{avoter}} = \frac{\sum_{u \in [n]} d_u x_u^{(0)}}{\sum_{u \in [n]} d_u}, \quad (2)$$

where $x_u^{(0)}$ and d_u are respectively the initial opinion and the degree of node u . Since $x_u^{(0)} \in \{0, 1\}$, the fixation probability ϕ^{avoter} is proportional to the volume of nodes initially holding opinion 1⁴.

In the analysis of the expected consensus time, instead, one cannot ignore the idle iterations, but since they occur with probability $1 - \alpha$ independently of other random choices, their effect is simply that of slowing down the standard asynchronous voter process by a factor $1/\alpha$. This intuitive argument can be formalized through a standard Markov chain coupling argument (see Appendix B).

Proposition 1. *In the unbiased asynchronous case, the fixation probability is the same as for the standard asynchronous voter model. The expected consensus time is T^{avoter}/α , where T^{avoter} is the expected consensus time of the standard asynchronous voter model.*

On the n -Clique Topology. As discussed in the introduction, a very natural question is: how large is the expected consensus time on the n -clique as a function of n ? Despite this question having been studied multiple times before, in the literature there are either diffusion approximations with unknown error [Ewens, 2012], or exact formulas involving multiple partial summations that are hard to interpret asymptotically [Glaz, 1979]. By further analyzing a result of Glaz [Glaz, 1979], we derive here an explicit formula with bounded error that is easy to interpret, which in fact agrees with the diffusion approximation up to lower order terms, thus also showing that at least in this case, the diffusion approximation (1) yields a correct estimate.

In the unbiased asynchronous model, the expected consensus time on the clique is the same as the mean absorption time of the underlying birth-death process, the state of which is summarized by the number of nodes holding opinion 1. Call $T_k(n)$ the expected consensus time when starting from a configuration with k nodes holding opinion 1, and the remaining $n - k$ holding opinion 0. We prove that $T_k(n) = \mathcal{O}(n^2/\alpha)$, more precisely:

Theorem 1. *If $\alpha_{01} = \alpha_{10} = \alpha$ (for some $\alpha > 0$), then for each $k = 1, \dots, n - 1$,*

$$T_k(n) = \frac{1}{\alpha} n^2 h(k/n) + \mathcal{O}(n/\alpha),$$

where $h(p) := -p \ln p - (1 - p) \ln(1 - p) \leq \ln 2$.

Proof. On an n -clique, the process is equivalent to a birth-and-death chain [Levin et al., 2009] on $n + 1$ states

⁴We note incidentally that the fixation probability can also be computed by suitably relating the asynchronous model to the transition matrix of the lazy random walk we discuss in Section 3.2. This connection is only mentioned here and made rigorous in the full paper version (Appendix A).

$0, 1, 2, \dots, n$ (representing the number of nodes with opinion, say, 1). Let us define the following quantities:

- $p_k = \alpha_{01} k(n - k)/n(n - 1)$ is the probability that the number of nodes holding opinion 1 increases from k to $k + 1$ when $0 \leq k < n$,
- $q_k = \alpha_{10} k(n - k)/n(n - 1)$ is the probability that the number of nodes holding opinion 1 decreases from k to $k - 1$ when $0 < k \leq n$.

Note that $p_k = q_k$ for all k due to the assumption $\alpha_{01} = \alpha_{10}$. Define the vector $T(n)$ as $T(n) = (T_1(n), \dots, T_{n-1}(n))^T$ and consider the matrix

$$B = \begin{pmatrix} p_1 + q_1 & -p_1 & 0 & \dots & 0 \\ -q_2 & p_2 + q_2 & -p_2 & \dots & 0 \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ 0 & \dots & 0 & -q_{n-1} & p_{n-1} + q_{n-1} \end{pmatrix}.$$

The matrix B is constructed so that $BT(n) = \mathbf{1}$, where $\mathbf{1}$ is the all-1 vector. This holds because of the recurrence

$$T_k(n) = 1 + (1 - p_k - q_k)T_k(n) + q_k T_{k-1}(n) + p_k T_{k+1}(n)$$

for the mean consensus times. Therefore, $T(n) = B^{-1}\mathbf{1}$. The matrix B can be explicitly inverted thanks to its tridiagonal structure; an explicit computation (see Appendix C) yields

$$T_k(n) = \frac{n-1}{\alpha} ((n-k)(H_{n-1} - H_{n-k}) + k(H_{n-1} - H_{k-1})),$$

where H_k is the k -th harmonic number, $H_k = \sum_{j=1}^k 1/j$. Recalling the asymptotic expansion $H_n = \ln n + \gamma + \mathcal{O}(1/n)$, where γ is the Euler-Mascheroni constant,

$$\begin{aligned} T_k(n) &= \frac{n-1}{\alpha} ((n-k)(H_n - H_{n-k}) + k(H_n - H_k)) + \mathcal{O}\left(\frac{n}{\alpha}\right) \\ &= \frac{n(n-1)}{\alpha} \left(\left(1 - \frac{k}{n}\right) \ln \frac{n}{n-k} + \frac{k}{n} \ln \frac{n}{k} \right) + \mathcal{O}(n/\alpha) \\ &= \frac{n^2}{\alpha} h(k/n) + \mathcal{O}(n/\alpha), \end{aligned}$$

where $h(p) = -p \ln p - (1 - p) \ln(1 - p)$, which is such that $0 \leq h(p) \leq \ln 2$ for every $p \in [0, 1]$. \square

3.2 Synchronous Variant

The analysis of the synchronous variant in the unbiased setting relies on the tight connection between the unbiased case of the opinion dynamics we consider and (lazy) random walks on networks.

Connections to Lazy Random Walks. We next provide an equivalent formulation of our model, which reveals an interesting and useful connection to lazy random walks. To this purpose, consider the following, alternative dynamics, in which the behavior of the generic node u at each iteration is the following:

- Node u independently tosses a coin with probability of “heads” equal to α ;
- If “heads”, u samples a neighbor v u.a.r. and copies v ’s opinion; otherwise u does nothing and keeps her opinion.

Let us call \mathcal{M}_1 the synchronous model described in Section 2 and \mathcal{M}_2 the dynamics described above. Then \mathcal{M}_1 and \mathcal{M}_2 are equivalent in the sense that, if they start from the same initial state, they generate the same probability distribution over all possible configurations of the system at any iteration t . Intuitively speaking this is true since in \mathcal{M}_1 each node first samples a neighbor and then it decides whether or not to copy its opinion according to the outcome of a coin toss, while in \mathcal{M}_2 each node first tosses a coin to decide whether or not to copy the opinion of one of the neighbors and then it samples the neighbor. Since the outcome of the coin toss and the choice of the neighbor are independent random variables, they produce the same distribution on the new opinion of the node when commuted. In the full paper version (Appendix D) we formalize the above equivalence and we prove it by appropriately coupling the two processes using an inductive argument.

Model \mathcal{M}_2 is interesting, since it describes a (lazy) voter model. As such (and as we explicitly show in the proof of Theorem 2), it is equivalent, in a probabilistic sense, to n lazy, coalescing random walks on the underlying network. This connection allows us to extrapolate the probability of consensus to a particular opinion and to adapt techniques that have been used to analyze the consensus time of the standard voter model [Hassin and Peleg, 2001; Aldous and Fill, 2002].

Theorem 2. *Assume model \mathcal{M}_2 starts in a configuration in which all nodes of a subset $W \subset V$ have opinion 1 and all other nodes have opinion 0. Let ϕ and T^{cons} denote fixation probability (of opinion 1) and time to consensus, resp. Then: (I) $\phi = (\sum_{u \in W} d_u) / (\sum_{u \in V} d_u)$, (II) $\mathbf{E}[T^{\text{cons}}] \leq \beta_n T^{\text{hit}}$, where T^{hit} is the maximum expected hitting time associated with the graph and $\beta_n = \mathcal{O}(1)$ when $\alpha \leq 1/2$, while $\beta_n = \ln n + 3$ when $\alpha > 1/2$.*

Sketch of the proof. We here only give a short idea of the proof and we defer a full-detailed proof to the full paper version (Appendix E).

The proof of (I) follows from the observation that, if we call $\mathbf{p}(t)$ the vector $\mathbf{p}(t) = (p_1(t), \dots, p_n(t))$ where $p_i(t)$ is the probability that node i has opinion 1 at round t conditional on the configuration at the previous round $\mathbf{x}^{(t-1)}$, then for every round t it holds that $\mathbf{E}[\mathbf{x}^{(t+1)} | \mathbf{x}^{(t)}] = \mathbf{p}(t+1) = P\mathbf{x}^{(t)}$ where P is the transition matrix of a lazy simple random walk on the underlying graph. Iterating the above equality we have that $\lim_{t \rightarrow \infty} \mathbf{E}[\mathbf{x}^{(t)} | \mathbf{x}^{(0)}] = \pi^\top \mathbf{x}^{(0)} \mathbf{1}$, where π is the stationary distribution of the random walk. Finally, the formula for ϕ follows from the fact that, for each node i , $\lim_{t \rightarrow \infty} \mathbf{E}[x_i^{(t)} | \mathbf{x}^{(0)}]$ equals the probability that node i ends up with opinion 1 and from the fact that the stationary probability of a simple random walk being on a node i is proportional to the degree of i .

The proof of (II) is an adaptation to our (discrete) case of the proof strategy for the continuous case described in [Aldous and Fill, 2002, Section 14.3.2]: we leverage on the relation between the convergence time of \mathcal{M}_2 and the maximum meeting time of two lazy random walks and, by using

an appropriate martingale, we show that the maximum meeting time is upper bounded by the maximum hitting time (see Appendix E). \square

4 The Biased Setting

Without loss of generality, in the rest of this section we assume $\alpha_{01} \neq \alpha_{10}$ and we let $r = \alpha_{01}/\alpha_{10}$. In general, the biased setting is considerably harder to address, since the connection between our model and lazy random walks no longer applies in this setting, nor does it seem easy to track the evolution of the expected behavior of the model in a way that is mathematically useful.

4.1 Asynchronous Variant

In the asynchronous case, we give a result for the fixation probability holding for regular graphs, thanks to an equivalence with the fixation probability for the n -clique (see Appendix F). We also bound the expected consensus time in the specific case of the n -clique. The asynchronous variant of the process on the n -clique is equivalent to a birth-and-death chain on $n+1$ states $0, 1, 2, \dots, n$ representing the number of nodes with opinion 1. Now, however, the transition probabilities will not be symmetric. In fact, the transition probabilities can be specified by $\{p_k, q_k, r_k\}_{k=0}^n$ where $p_k + q_k + r_k = 1$, and:

- p_k is the probability that the number of nodes holding opinion 1 increases from k to $k+1$ when $0 \leq k < n$,
- q_k is the probability that the number of nodes holding opinion 1 decreases from k to $k-1$ when $0 < k \leq n$,
- r_k is the probability that the number of nodes holding opinion 1 remains k when $0 \leq k \leq n$.

Due to our definition of the opinion dynamics, we have $p_0 = q_n = 0$ and, for $0 < k < n$,

$$\begin{aligned} p_k &= \left(1 - \frac{k}{n}\right) \cdot \frac{k}{n-1} \cdot \alpha_{01} = \alpha_{01} \frac{k(n-k)}{n(n-1)}, \\ q_k &= \frac{k}{n} \cdot \frac{n-k}{n-1} \cdot \alpha_{10} = \alpha_{10} \frac{k(n-k)}{n(n-1)}, \\ r_k &= 1 - (\alpha_{01} + \alpha_{10}) \frac{k(n-k)}{n(n-1)}. \end{aligned}$$

Theorem 3. *Let $r = \alpha_{01}/\alpha_{10}$ and let ϕ_k be the fixation probability of opinion 1 on a regular n -nodes graph starting from a state in which k nodes hold opinion 1. Then, for $r \notin \{0, 1\}$,*

$$\phi_k = \frac{1 - r^{-k}}{1 - r^{-n}}. \quad (3)$$

Proof. Thanks to the equivalence of the fixation probability between a regular n -nodes graph and an n -clique (see Appendix F), and by the analysis of a general birth-death process (see for example [Nowak, 2006, Section 6.2]), we get

$$\phi_k = \frac{1 + \sum_{i=1}^{k-1} \prod_{j=1}^i \gamma_j}{1 + \sum_{i=1}^{n-1} \prod_{j=1}^i \gamma_j},$$

where $\gamma_j = q_j/p_j = \alpha_{10}/\alpha_{01} = 1/r$ for all j . Hence

$$\phi_k = \frac{1 + \sum_{i=1}^{k-1} r^{-i}}{1 + \sum_{i=1}^{n-1} r^{-i}} = \frac{1 + \frac{1-r^{-k}}{1-r^{-1}} - 1}{1 + \frac{1-r^{-n}}{1-r^{-1}} - 1} = \frac{1 - r^{-k}}{1 - r^{-n}}.$$

Note that when $r \rightarrow 1$, we can evaluate ϕ_k by applying L'Hôpital's rule to (3) and get $\phi_k = k/n$, which is consistent with the results for the unbiased setting. When $r = 0$, nodes can only switch from opinion 1 to opinion 0, so clearly $\phi_k = 0$; this is also consistent with (3) when $r \rightarrow 0$. \square

It is interesting to note that the expression from (3) coincides with the fixation probability in the standard Moran process [Nowak, 2006, Chapter 6] when mutants (say, nodes with opinion 1) have a *relative fitness* equal to α_{01}/α_{10} , and in the initial configuration there are k mutants out of n nodes. In other words, on an n -nodes regular graph the ratio α_{01}/α_{10} can be interpreted as a *fitness* of sorts, even though there is no notion of fitness or selection built in our model (recall that nodes are activated uniformly at random).

For the n -clique we are also able to bound the expected consensus time. While the logic of the proof is similar to the one of Theorem 1, the proof itself is considerably more involved in the asymmetric setting, leading to qualitatively different results—namely, an $O(n \log n)$ instead of an $O(n^2)$ worst case bound. The reader is deferred to the full paper version (Appendix G) for full details.

Theorem 4. *If $T_k(n)$ is the expected consensus time in the n -clique when starting from a state with k nodes holding opinion 1, and α_{01}, α_{10} are constants, then for each $k = 1, \dots, n-1$,*

$$T_k(n) = O(n \log n),$$

and for some values of k the above bound is tight.

4.2 Synchronous Variant

In order to bound the fixation probabilities, denote by $\mathbf{x}^{(t)} \in \{0, 1\}^n$ the state of the system at time t . Conditioned on the state vector $\mathbf{x}^{(t)}$, the probability that $x_u^{(t+1)} = 1$ can be expressed as follows:

$$\mathbf{P} \left(x_u^{(t+1)} = 1 \mid \mathbf{x}^{(t)} \right) = \begin{cases} 1 - \alpha_{10} \left(1 - \frac{\sum_{v \in V} a_{uv} x_v^{(t)}}{d_u} \right) & \text{if } x_u^{(t)} = 1 \\ \alpha_{01} \frac{\sum_{v \in V} a_{uv} x_v^{(t)}}{d_u} & \text{if } x_u^{(t)} = 0. \end{cases}$$

This follows since the probability that node u samples a neighbor with opinion 1 is $\sum_v a_{uv} x_v^{(t)} / d_u$ and:

- when $x_u^{(t)} = 1$, then $x_u^{(t+1)} = 1$ iff either u samples a neighbor with opinion 1, or u samples a neighbor with opinion 0 and does not accept its opinion (these two events are disjoint);
- when $x_u^{(t)} = 0$, then $x_u^{(t+1)} = 1$ iff u samples a neighbor with opinion 1 and accepts its opinion.

We did not exploit the graph topology so far. In the case of the n -clique (with loops, to simplify some expressions), let $k^{(t)}$ be the number of nodes with opinion 1 at time t . Specializing the formulas derived above we get

$$\mathbf{P} \left(x_u^{(t+1)} = 1 \mid \mathbf{x}^{(t)} \right) = \begin{cases} 1 - \alpha_{10} \left(1 - \frac{k^{(t)}}{n} \right) & \text{if } x_u^{(t)} = 1 \\ \alpha_{01} \frac{k^{(t)}}{n} & \text{if } x_u^{(t)} = 0. \end{cases}$$

Note that the expression above depends only on $k^{(t)}$ and $x_u^{(t)}$, and not on the entire state $\mathbf{x}^{(t)}$. The process is thus equivalent to sampling, at each step t , $k^{(t)}$ Bernoulli random variables (r.v.) with parameter $\beta_k := 1 - \alpha_{10}(1 - k^{(t)}/n)$, and $n - k^{(t)}$ Bernoulli r.v. with parameter $\gamma_k := \alpha_{01}k^{(t)}/n$. Collectively, the outcomes of these r.v. constitute the new state $\mathbf{x}^{(t+1)}$.

Then,

$$\mathbf{E} \left[k^{(t+1)} \mid k^{(t)} \right] = (n - k^{(t)})\alpha_{01} \frac{k^{(t)}}{n} + k^{(t)} \left(1 - \alpha_{10} \left(1 - \frac{k^{(t)}}{n} \right) \right)$$

which, posing $y^{(t)} = k^{(t)}/n$, can be written as

$$\mathbf{E} \left[y^{(t+1)} \mid y^{(t)} \right] = y^{(t)} + (\alpha_{01} - \alpha_{10})y^{(t)}(1 - y^{(t)}). \quad (4)$$

Proposition 2. *Assume $\alpha_{01} \leq \alpha_{10}$. Then the fixation probability of opinion 0 is at least the fraction of agents holding opinion 0.*

Proof. Under the assumption $\alpha_{01} \leq \alpha_{10}$,

$$\begin{aligned} \mathbf{E} \left[y^{(t+1)} \right] &= \mathbf{E} \left[\mathbf{E} \left[y^{(t+1)} \mid y^{(t)} \right] \right] \\ &= \mathbf{E} \left[y^{(t)} \right] + (\alpha_{01} - \alpha_{10})\mathbf{E} \left[y^{(t)}(1 - y^{(t)}) \right] \leq \mathbf{E} \left[y^{(t)} \right]. \end{aligned}$$

Hence, the succession $(\mathbf{E} [y^{(t)}])_t$ is monotone and bounded and attains a limit. This limit must coincide with the fixation probability, because $y^{(t)}$ converges in distribution to a Bernoulli random variable $y^{(\infty)}$ and $\mathbf{E} [y^{(\infty)}] = \mathbf{P} (y^{(\infty)} = 1) = \mathbf{P} (\exists t : y^{(t)} = 1)$ equals the fixation probability. Since $\mathbf{E} [y^{(0)}] = y^{(0)} = k^{(0)}/n$, the fixation probability of opinion 1 must be at most $k^{(0)}/n$, so that of opinion 0 is at least $1 - k^{(0)}/n$. \square

Regarding the expected consensus time, we show the following by using the technique of *drift analysis* [Lengler and Steger, 2018].

Theorem 5. *If $T_k(n)$ is the expected consensus time in the n -clique when starting from a configuration with k nodes holding opinion 1, and $\alpha_{01} = \alpha_{10} - \epsilon$, then*

$$T_k(n) \leq \frac{nk}{\epsilon(n-1)}.$$

In particular, $T_k(n) \leq \min(2k/\epsilon, n/\epsilon)$ for each $k = 1, 2, \dots, n-1$.

Proof. We adapt a proof of [Lengler and Steger, 2018, Theorem 2.1] to our setting, since their result is not suitable for systems with more than one absorbing state. In the remainder, we assume $\alpha_{01} < \alpha_{10}$, we let $\epsilon = \alpha_{10} - \alpha_{01}$, and we let $z^{(t)} = k^{(t)}/n$, i.e., $z^{(t)}$ is the fraction of agents with opinion 1 at time t . We begin by defining the following stopping time:

$$T := \inf \{ t \geq 0 : z^{(t)} \in \{0, 1\} \}.$$

This definition is akin to the one given in [Lengler and Steger, 2018, Theorem 2.1], but it accounts for the presence of two absorbing states in the Markov chain defined by $z^{(t)}$. Moreover, $z^{(t)} = z^{(t-1)}$ for every $t > T$, since

$z^{(t)}$ does not change after absorption (regardless of the absorbing state). We next note that for every t , $z^{(t)} \in \mathcal{S} = \{0, \frac{1}{n}, \dots, 1 - \frac{1}{n}, 1\}$. Moreover, for every $s \in \mathcal{S}$ we have $\mathbf{E}[z^{(t+1)} | z^{(t)} = s] = s - \epsilon s(1 - s)$, whence:

$$\mathbf{E}[z^{(t)} - z^{(t+1)} | z^{(t)} = s] = \epsilon s(1 - s),$$

with the last quantity at least $\epsilon \frac{1}{n} (1 - \frac{1}{n})$ for $s \in \mathcal{S} \setminus \{0, 1\}$. Next:

$$\begin{aligned} & \mathbf{E}[z^{(t+1)} | T > t] \\ &= \sum_{s=1}^{n-1} \mathbf{E}[z^{(t+1)} | z^{(t)} = \frac{s}{n}] \cdot \mathbf{P}\left(z^{(t)} = \frac{s}{n} | T > t\right), \end{aligned}$$

where the first equality follows since, for $s \notin \{0, 1\}$, $z^{(t)} = s$ implies $T > t$. Similarly to the proof of [Lengler and Steger, 2018, Theorem 2.1], the equality above implies

$$\mathbf{E}[z^{(t)} - z^{(t+1)} | T > t] \geq \epsilon \frac{1}{n} \left(1 - \frac{1}{n}\right). \quad (5)$$

We let $\delta = \epsilon \frac{1}{n} (1 - \frac{1}{n})$ for conciseness. We next have:

$$\begin{aligned} & \mathbf{E}[z^{(t)}] \stackrel{(a)}{=} \mathbf{E}[z^{(t)} | T > t] \cdot \mathbf{P}(T > t) + \\ & \quad + \mathbf{P}(z^{(t)} = 1 | T \leq t) \cdot \mathbf{P}(T \leq t) \end{aligned}$$

and therefore

$$\begin{aligned} & \mathbf{E}[z^{(t+1)}] \stackrel{(a)}{=} \mathbf{E}[z^{(t+1)} | T > t] \cdot \mathbf{P}(T > t) + \\ & \quad + \mathbf{P}(z^{(t+1)} = 1 | T \leq t) \cdot \mathbf{P}(T \leq t) \\ & \stackrel{(b)}{\leq} \left(\mathbf{E}[z^{(t)} | T > t] - \delta\right) \cdot \mathbf{P}(T > t) + \\ & \quad + \mathbf{P}(z^{(t+1)} = 1 | T \leq t) \cdot \mathbf{P}(T \leq t) \\ & \stackrel{(c)}{=} \mathbf{E}[z^{(t)}] - \mathbf{P}(z^{(t)} = 1 | T \leq t) \cdot \mathbf{P}(T \leq t) + \\ & \quad + \mathbf{P}(z^{(t+1)} = 1 | T \leq t) \cdot \mathbf{P}(T \leq t) - \delta \cdot \mathbf{P}(T > t). \end{aligned}$$

In the derivations above, (a) simply follows from the law of total probability, considering that $T \leq t$ implies $z^{(t)} \in \{0, 1\}$, (b) follows from (5), while (c) follows by replacing the equation of $\mathbf{E}[z^{(t)}]$ into the last step of the derivation. Next, we note that

$$\mathbf{P}(z^{(t+1)} = 1 | T \leq t) = \mathbf{P}(z^{(t)} = 1 | T \leq t)$$

by definition of the $z^{(t)}$, whence we obtain:

$$\delta \cdot \mathbf{P}(T > t) \leq \mathbf{E}[z^{(t)}] - \mathbf{E}[z^{(t+1)}]. \quad (6)$$

Now, observe that (6) is exactly [Lengler and Steger, 2018, (2.4) in Theorem 2.1]. From this point, the proof proceeds exactly as in [Lengler and Steger, 2018, (2.4) in Theorem 2.1], so that we finally have:

$$\mathbf{E}[T] \leq \frac{z^{(0)}}{\delta} = \frac{nk}{\epsilon(n-1)},$$

if at time $t = 0$ we have k agents with opinion 1. \square

5 Conclusions and Outlook

Natural directions for future work include considering more opinions and general topologies.

More Opinions. The case of more opinions presents no major challenges in the unbiased case, both in its asynchronous and synchronous variants, something we did not discuss for the sake of space. In this case, one can simply focus on one opinion at a time, collapsing the remaining opinions into an “other” class. Proceeding this way, it is easy to extend the results we presented in Section 3 to the general case: for $k > 2$ opinion, the fixation probability for opinion i is $\frac{\sum_{u \in W} d_u}{\sum_u d_u}$, where W is the subset of nodes with opinion i in the initial configuration. The biased case is considerably harder and the technical barriers are twofold: one is the general difficulty of characterizing the expected change of the global state in the biased setting even in the case of 2 opinions (see next paragraph). The other is the possible presence of rock-paper-scissors like dynamics that may arise depending on the distribution of the opinion biases.

General Topologies. As also suggested by previous work, albeit for different models [Montanari and Saberi, 2010; Anagnostopoulos *et al.*, 2022; Lesfari *et al.*, 2022], we believe the biased case might give rise to diverse and possibly counterintuitive behaviors. In general, a crucial technical challenge is characterizing the evolution of the global state across consecutive steps, since this in general depends on the current configuration in a way that is highly topology-dependent and hard to analyze. Some recent results [Schoenebeck and Yu, 2018; Shimizu and Shiraga, 2020] proposed techniques relying on variants of the expander mixing lemma to investigate quasi-majority dynamics on expanders. Unfortunately, these techniques do not obviously extend to the biased voter models we consider. Indeed and interestingly, the class of dynamics these techniques apply to does not even include the standard voter model as a special case.

In general, we believe that extending and/or improving our results for the biased setting might require refining important techniques, such as those of [Schoenebeck and Yu, 2018; Shimizu and Shiraga, 2020] or the ones discussed in [Lengler and Steger, 2018].

Acknowledgements

- Partially supported by the ERC Advanced Grant 788893 AMDROMA “Algorithmic and Mechanism Design Research in Online Markets”, the EC H2020RIA project “SoBig-Data++” (871042), and the MIUR PRIN project ALGADIMAR “Algorithms, Games, and Digital Markets.
- Work of the second author was carried out in association with Istituto di Analisi dei Sistemi ed Informatica, Consiglio Nazionale delle Ricerche, Italy.
- Supported by the Austrian Science Fund (FWF): P 32863-N.
- 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 947702).

References

- [Aldous and Fill, 2002] David Aldous and James Fill. Reversible markov chains and random walks on graphs. <https://www.stat.berkeley.edu/~aldous/RWG/book.pdf>, 2002. Accessed: 2023-05-22.
- [Anagnostopoulos *et al.*, 2022] Aris Anagnostopoulos, Luca Becchetti, Emilio Cruciani, Francesco Pasquale, and Sara Rizzo. Biased opinion dynamics: when the devil is in the details. *Information Sciences*, 593:49–63, 2022.
- [Assaf and Mobilia, 2012] Michael Assaf and Mauro Mobilia. Metastability and anomalous fixation in evolutionary games on scale-free networks. *Physical review letters*, 109(18):188701, 2012.
- [Baxter *et al.*, 2008] G. J. Baxter, R. A. Blythe, and A. J. McKane. Fixation and consensus times on a network: A unified approach. *Physical Review Letters*, 101(25):258701, December 2008.
- [Berenbrink *et al.*, 2016] Petra Berenbrink, George Giakkoupis, Anne-Marie Kermarrec, and Frederik Mallmann-Trenn. Bounds on the voter model in dynamic networks. In *43rd International Colloquium on Automata, Languages, and Programming, ICALP 2016*, volume 55 of *LIPICs*, pages 146:1–146:15, 2016.
- [Cooper *et al.*, 2018] Colin Cooper, Martin Dyer, Alan Frieze, and Nicolás Rivera. Discordant voting processes on finite graphs. *SIAM Journal on Discrete Mathematics*, 32(4):2398–2420, 2018.
- [Cruciani *et al.*, 2021] Emilio Cruciani, Hlafo Alfie Mimun, Matteo Quattropiani, and Sara Rizzo. Phase transitions of the k-majority dynamics in a biased communication model. In *ICDCN '21: International Conference on Distributed Computing and Networking*, pages 146–155. ACM, 2021.
- [Díaz *et al.*, 2016] Josep Díaz, Leslie Ann Goldberg, David Richerby, and Maria Serna. Absorption time of the Moran process. *Random Structures & Algorithms*, 49(1):137–159, 2016.
- [Donnelly and Welsh, 1983] Peter Donnelly and Dominic Welsh. Finite particle systems and infection models. *Mathematical Proceedings of the Cambridge Philosophical Society*, 94(1):167–182, July 1983. Publisher: Cambridge University Press.
- [Durocher *et al.*, 2022] Loke Durocher, Panagiotis Karras, Andreas Pavlogiannis, and Josef Tkadlec. Invasion dynamics in the biased voter process. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence, IJCAI 2022*, pages 265–271. ijcai.org, 2022.
- [Ewens, 2012] Warren Ewens. *Mathematical Population Genetics I: Theoretical Introduction*. Springer, 2012.
- [Glaz, 1979] Joseph Glaz. Probabilities and moments for absorption in finite homogeneous birth-death processes. *Biometrics*, 35(4):813–816, 1979.
- [Hassin and Peleg, 2001] Yehuda Hassin and David Peleg. Distributed probabilistic polling and applications to proportionate agreement. *Inf. Comput.*, 171(2):248–268, 2001.
- [Hindersin and Traulsen, 2014] Laura Hindersin and Arne Traulsen. Counterintuitive properties of the fixation time in network-structured populations. *Journal of The Royal Society Interface*, 11(99):20140606, 2014.
- [Kimura, 1962] Motoo Kimura. On the probability of fixation of mutant genes in a population. *Genetics*, 47(6):713–719, 1962.
- [Lengler and Steger, 2018] Johannes Lengler and Angelika Steger. Drift analysis and evolutionary algorithms revisited. *Combinatorics, Probability and Computing*, 27(4):643–666, July 2018.
- [Lesfari *et al.*, 2022] Hicham Lesfari, Frédéric Giroire, and Stéphane Pérennes. Biased majority opinion dynamics: Exploiting graph k-domination. In Luc De Raedt, editor, *Proceedings of the 31st International Joint Conference on Artificial Intelligence, IJCAI 2022*, pages 377–383. ijcai.org, 2022.
- [Levin *et al.*, 2009] David A. Levin, Yuval Peres, and Elizabeth L. Wilmer. *Markov Chains and Mixing Times*. American Mathematical Society, 2009.
- [Lieberman *et al.*, 2005] Erez Lieberman, Cristoph Hauert, and Martin A. Nowak. Evolutionary dynamics on graphs. *Nature*, 433:312–316, 2005.
- [Liggett, 1985] Thomas M. Liggett. *Interacting Particle Systems*. Springer New York, 1985.
- [Montanari and Saberi, 2010] Andrea Montanari and Amin Saberi. The spread of innovations in social networks. *Proceedings of the National Academy of Sciences*, 107(47):20196–20201, 2010.
- [Moran, 1958] Patrick A. P. Moran. Random processes in genetics. *Mathematical Proceedings of the Cambridge Philosophical Society*, 54(1):60–71, January 1958.
- [Nowak, 2006] Martin A. Nowak. *Evolutionary Dynamics*. Harvard University Press, 2006.
- [Peralta *et al.*, 2021] Antonio F Peralta, Matteo Neri, János Kertész, and Gerardo Iñiguez. Effect of algorithmic bias and network structure on coexistence, consensus, and polarization of opinions. *Physical Review E*, 104(4):044312, 2021.
- [Schoenebeck and Yu, 2018] Grant Schoenebeck and Fang-Yi Yu. Consensus of interacting particle systems on erdős-rényi graphs. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1945–1964. SIAM, 2018.
- [Shimizu and Shiraga, 2020] Nobutaka Shimizu and Takeharu Shiraga. Quasi-majority functional voting on expander graphs. In *47th International Colloquium on Automata, Languages, and Programming, ICALP 2020*, volume 168 of *LIPICs*, pages 97:1–97:19. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.
- [Sood *et al.*, 2008] Vishal Sood, Tibor Antal, and Sidney Redner. Voter models on heterogeneous networks. *Physical Review E*, 77(4):041121, April 2008.